arXiv:1704.02381v1 [stat.ME] 7 Apr 2017

Adaptive estimation of the rank of the coefficient matrix in high dimensional multivariate response regression models

Xin Bing[1] and Marten Wegkamp[1,2]

[1]Department of Statistical Science, Cornell University
[2]Department of Mathematics, Cornell University

Abstract

We consider the multivariate response regression problem with a regression coefficient matrix of low, unknown rank. In this setting, we analyze a new criterion for selecting the optimal reduced rank. This criterion differs notably from the one proposed in Bunea, She and Wegkamp [7] in that it does not require estimation of the unknown variance of the noise, nor depends on a delicate choice of a tuning parameter. We develop an iterative, fully data-driven procedure, that adapts to the optimal signal to noise ratio. This procedure finds the true rank in a few steps with overwhelming probability. At each step, our estimate increases, while at the same time it does not exceed the true rank. Our finite sample results hold for any sample size and any dimension, even when the number of responses and of covariates grow much faster than the number of observations. We perform an extensive simulation study that confirms our theoretical findings. The new method performs better and more stable than that in [7] in both low- and high-dimensional settings.

1 Introduction

1.1 Background

We study the multivariate response regression model

$$Y = XA + E \in \mathbb{R}^{n\times m}$$

with $X \in \mathbb{R}^{n\times p}$ of $\text{rank}(X) = q$ and $A \in \mathbb{R}^{p\times m}$ of unknown $\text{rank}(A) = r$. We assume that the entries $[E]_{ij}$ of E are i.i.d. $N(0, \sigma^2)$ distributed with $\sigma^2 < \infty$.

Standard least squares estimation is tantamount to regressing each response on the predictors separately, thus ignoring the multivariate nature of the possibly correlated responses. In large dimensional settings (m and p are large relative to the sample size n), it is desirable to achieve a dimension reduction in the coefficient matrix A. One popular way of achieving this goal, is to find a common subset of $s \leq p$ covariates that are relevant for prediction, using penalized least squares with a ℓ_1/ℓ_2 (group lasso) type penalty on the regression coefficients, see, for instance, [8, 9, 23, 27, 41] to recover the support of the set of s rows for which A is non-zero.

Reduced rank regression is a different approach to achieve necessary dimension reduction. The main premise is that A has low rank, so that we can write $A = A_1A_2$ for a $p \times r$ matrix A_1 and a $r \times m$ matrix A_2. Then only few linear combinations $X^* = XA_1$ of X are needed to explain the variation of Y. Izenman [19] coined the term *reduced-rank regression* for this class of models, but its history dates back to Anderson [1]. There are many works on this topic in the classical setting of *fixed* dimensions m and p, and sample size $n \to \infty$, see [28, 30, 31] and more recently, [2, 3]. A comprehensive overview on reduced rank regression is given by [29]. Only recently, the high-dimensional case has been discussed: [7, 8, 17, 18, 26, 32].

The main topic of this paper is the estimation of the unknown rank. Determination of the rank of the coefficient matrix is the first key step for the estimation of A. For known rank r, Anderson [2] derives the asymptotic distribution of the reduced rank regression coefficient matrix estimator $\widehat{A}_r$ in the asymptotic setting with m, p fixed and $n \to \infty$. The estimator $\widehat{A}_k$ is the matrix corresponding to minimizing the squared Frobenius norm $\|Y - XB\|^2$ over all $p \times n$ matrices B of rank k and has a closed form, due to the Eckart-Young theorem[14, 34]. It is crucial to have the true rank $k = r$ for obtaining a good fit for both $\|XA - X\widehat{A}_k\|^2$ and $\|A - \widehat{A}_k\|^2$. Figure 1 depicts the role of r. In general, the rank r is unknown *a priori*. The classical approach to estimate the rank r uses the likelihood ratio test, see [1]. An elementary calculation shows that this statistic coincides with Bartlett's test statistic as a consequence of the relation between reduced rank regression and canonical correlation analysis, see [1, 33]. Our main goal in this study is to develop a non-asymptotic method to estimate r that is

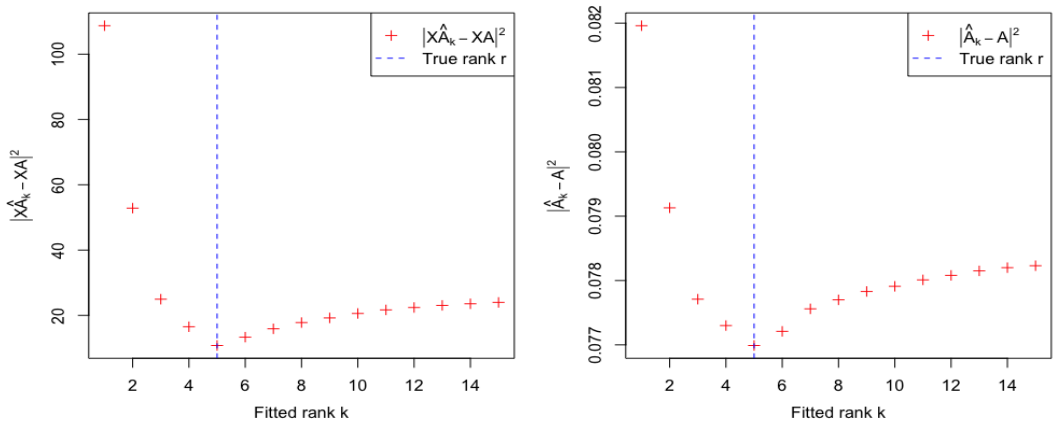

Figure 1: Plots of $40 \cdot \|X\widehat{A}_k - XA\|^2/(nm)$ and $40 \cdot \|\widehat{A}_k - A\|^2/(pm)$ versus k in the dimension setting $n = 50, m = 30, p = 150, q = 15, r = 5, \eta = 0.1, b_0 = 0.02$ and $\sigma^2 = 2$.

easy to compute, adaptively from the data, and valid for any values of m, n and p, especially when the number of predictors p is large. The resulting estimator of A can then be used to construct a possibly much smaller number of new transformed predictors, or the most important canonical variables based on the original X and Y, see [20, Chapter 6] for a historical account. Under weak assumptions on the signal, our estimate of r can be shown to be equal to r with overwhelming probability, to wit, $1 - \exp(-\theta_1 mn) - \exp(-\theta_2(m+q))$, for some positive, finite constants θ_1, θ_2, so that the selection error is small compared to the overall error in estimating A.

1.2 Recent developments

Bunea, She and Wegkamp [7, 8] propose

$$\min_A \left\{ \|Y - XA\|^2 + \mu \cdot \mathrm{rank}(A) \right\} \tag{1}$$

and recommended the choice $\mu = C(\sqrt{m} + \sqrt{q})^2\sigma^2$ with constant $C > 1$ for the tuning parameter μ. Here $\|\cdot\|$ is the Frobenius norm. In particular, [7] established a convenient closed form for the $\mathrm{rank}(\widehat{A}) = \widehat{r}$ of the matrix $\widehat{A}$ that minimizes criterion (1). They gave sufficient conditions on the level of the smallest non-zero singular value of XA to guarantee that $\widehat{r}$ consistently estimates r. The disadvantage of this method is that a value for σ^2, in addition to the tuning parameter C, is required for μ. Bunea, She and Wegkamp [7] proposed

to use the unbiased estimator

$$\widetilde{\sigma}^2 := \frac{\|Y - PY\|^2}{nm - qm} \tag{2}$$

based on the projection PY of Y onto the range space of X. However, this becomes problematic when $(n-q)m$ is not large enough, or even infeasible when $n = q$.
Giraud [17] introduces another estimation scheme, that does not require estimation of σ^2. Unfortunately, a closed form for the minimizer as in [7] is lacking, and rank consistency in fact fails, as our simulations reveal in Section 5.7. Both procedures are rather sensitive to the choices of their respective tuning parameters involved.

1.3 Proposed Research

This paper studies a third criterion,

$$\widehat{\sigma}_k^2 := \frac{\|Y - (PY)_k\|^2}{nm - \lambda k}. \tag{3}$$

Here $(PY)_k = \sum_{j=1}^{k} d_j(PY) u_j v_j^T$ is the truncated singular value decomposition of PY based on the (decreasing) singular values $d_j(PY)$ of the projection PY and their corresponding singular vectors u_j, v_j. It was suggested by the second author in a private conversation to Christophe Giraud, see [18]. The range over which we minimize (3) is $\{0, 1, \ldots, K\}$ with $K = K_\lambda := \left(\lfloor \frac{nm-1}{\lambda} \rfloor \wedge m \wedge q \right)$ to avoid a non-positive denominator. While we improve upon the oracle inequality of the minimizer of this criterion given in Giraud [18] in the appendix, the purpose of this paper is to show that this new criterion produces a consistent estimator of the rank. It turns out that the choice of the optimal tuning parameter λ involves a delicate trade-off. On the one hand, λ should be large enough to prevent overfitting, that is, prevent selecting a rank that is larger than the true rank r. On the other hand, if one takes λ too large, the selected rank will typically be smaller than r as the procedure will not be able to distinguish the unknown singular values $d_j(XA)$ from the noise for $j > s$, for some $s = s(\lambda) < r$. To effectively deal with this situation, we refine our initial procedure using our new criterion (3) by an iterative procedure in Section 4 that provenly finds the optimal value of λ and consequently of the estimate of r. This method does not require any data-splitting and our simulations show that it is very stable. To our knowledge, it is rare feat to have a feasible algorithm that finds the optimal tuning parameter in a fully data-driven way, without data-splitting, and with mathematically proven guarantees.

The paper is organized as follows. Section 2 shows that the minimizer of (3) has a closed form. The main results are discussed in Sections 3 and 4. It obtains rank consistency in case of no signal ($XA = 0$) and in case of sufficient signal. For the latter, we develop a key notion of signal to noise ratio that is required for rank consistency. A sufficient, easily interpretable condition will be presented that corresponds to a computable value (estimate) of the tuning parameter λ. We develop in Section 4 an iterative, fully automated procedure, which has a guaranteed recovery of the true rank (with overwhelming probability) under increasingly milder conditions on the signal. The first step uses the potentially suboptimal estimate $\widehat{k}_0$ developed in Section 3, but which is less than r, with overwhelming probability. This value $\widehat{k}_0$ is used to update the tuning parameter λ. Then we minimize (3) again, and obtain a new estimate $\widehat{k}_1$, which in turn is used to update λ. The procedure produces each time a smaller λ, thereby selecting a larger rank k than the previous one, whilst each time we can guarantee that the selected ranks don't exceed the true rank r. This is a major mathematical challenge in our proof, which relies on highly non-trivial monotonicity arguments. The procedure stops when the selected rank does not change after an iteration. Our results hold with high probability (exponential in mn and $m + q$) which translates into extremely high accuracy (of our estimates) if the signal to noise ratio is sufficiently large. A large simulation study is reported in Section 5, which confirms our theoretical findings, and shows that our method dominates the methods proposed in [7, 17]. Most proofs are deferred to the appendix, while technical proofs can be found in the appendix. The implementation and simulation code can be found in https://github.com/bingx1990/Self-Tuning-Procedure-for-Reduced-Rank-Regression.git.

1.4 Notation

For any matrix A, we will use $[A]_{k\ell}$ to denote the k,ℓth element of A (i.e., the entry on the kth row and ℓth column of A), and we write $d_1(A) \geq d_2(A) \geq \cdots$ to denote its ordered singular values.

The Frobenius inner product $\langle \cdot, \cdot \rangle$ on the space of matrices is defined as $\langle A, B \rangle = \mathrm{tr}(A^T B)$ for commensurate matrices A, B. The corresponding Frobenius norm is denoted by $\|\cdot\|^2$, and we recall that $\|A\|^2 = \mathrm{tr}(A^t A) = \sum_j d_j^2(A)$ for any matrix A. Moreover, it is known (Eckhart-Young [14], Schmidt [34], see also the review by Stewart [35]) that minimizing $\|A - B\|^2$ over B with $r(B) \leq r$ is achieved for $B = (A)_r = UD_rV^T$ based on the singular value

decomposition of $A = UDV^T$ where D_r denotes the diagonal matrix with $[D]_{ii} = d_i(A)$ for $i = 1, \ldots, r$. Hence, $\min_{B:\ r(B)=r} \|A - B\|^2 = \sum_{j>r} d_j^2(A)$.
For general $m \times n$ matrices A and B, Weyl's theorem [39] implies that $d_{i+j-1}(A + B) \leq d_i(A) + d_j(B)$ for $1 \leq i, j, \leq q$ and $i + j \leq q + 1$ with $q = \min\{m, n\}$.
For other norms on matrices, we use $\|\cdot\|_2$ to denote the operator norm and $\|\cdot\|_*$ the nuclear norm (i.e., the sum of singular values). We have the inequalities $< A, B >= \mathrm{tr}(A^T B) \leq \|A\|_2 \|B\|_*$ and $\|A\|_* \leq \sqrt{\mathrm{rank}(A)}\|A\|$.
We denote the projection matrix onto the column space of X by P and we write $q := \mathrm{rank}(X) = r(P)$ and $N := r(PY) \leq q \wedge m$. We set $\widehat{\sigma}_0^2 := \|Y\|^2/(nm)$, by defining $(PY)_0 := 0$. Finally, we define $\rho := nm/(nm - \lambda K)$ and $\widehat{\sigma}^2 := \|E\|^2/(nm)$.
Throughout the paper, we use A to denote the true coefficient matrix and r to denote its true rank.

2 Properties of the minimizer of the new criterion

At first glance, it seems difficult to describe the minimizer $\widehat{k}$ of $k \mapsto \widehat{\sigma}_k^2$ because *both* the numerator and denominator in $\widehat{\sigma}_k^2$ are decreasing in k. However, it turns out that there is a unique minimizer with a neat explicit formula. First we characterize the comparison between $\widehat{\sigma}_i$ and $\widehat{\sigma}_j$ for $i \neq j$.

Proposition 1. *Let $i, j \in \{0, 1, \ldots, K\}$ with $i < j$. Then*

$$\widehat{\sigma}_j^2 \leq \widehat{\sigma}_i^2 \iff \frac{1}{j-i} \sum_{k=i+1}^{j} d_k^2(PY) \geq \lambda \widehat{\sigma}_j^2. \tag{4}$$

In particular,

$$\widehat{\sigma}_j^2 \leq \widehat{\sigma}_{j-1}^2 \iff d_j^2(PY) \geq \lambda \widehat{\sigma}_j^2, \tag{5}$$

and

$$d_j^2(PY) \leq \lambda \widehat{\sigma}_j^2 \iff d_j^2(PY) \leq \lambda \widehat{\sigma}_{j-1}^2. \tag{6}$$

This result and the monotonicity of the singular values $d_1(PY) \geq d_2(PY) \geq \cdots$ readily yield the following statement.

Proposition 2. *Let $k \in \{1, \dots, K\}$. Then*

$$\widehat{\sigma}_k^2 \le \widehat{\sigma}_{k-1}^2 \iff \widehat{\sigma}_k^2 \le \widehat{\sigma}_\ell^2 \text{ for all } \ell \le k-1. \tag{7}$$

$$\widehat{\sigma}_k^2 \ge \widehat{\sigma}_{k-1}^2 \iff \widehat{\sigma}_\ell^2 \ge \widehat{\sigma}_{k-1}^2 \text{ for all } \ell > k-1. \tag{8}$$

It is clear that if $\widehat{\sigma}_1^2 \ge \widehat{\sigma}_0^2$, then $k = 0$ minimizes $\widehat{\sigma}_k^2$. Likewise, if $\widehat{\sigma}_K^2 \le \widehat{\sigma}_{K-1}^2$, then $k = K$ minimizes the criterion $\widehat{\sigma}_k^2$. After a little reflexion, we see that $\widehat{\sigma}_k^2$ is minimized at the last k for which $d_k^2(PY) \ge \lambda \widehat{\sigma}_k^2$ holds. That is,

$$\widehat{k} = \max\left\{0 \le k \le K : d_j^2(PY) \ge \lambda \widehat{\sigma}_j^2 \text{ for all } j \le k \text{ and } d_{k+1}^2(PY) < \lambda \widehat{\sigma}_{k+1}^2\right\} \tag{9}$$

minimizes $\widehat{\sigma}_k^2$ with the convention that the maximum of the empty set is 0. Properties (7) and (8) ensure that $d_j^2(PY) \ge \lambda \widehat{\sigma}_j^2$ must hold automatically for all $j \le k$ as well as $d_\ell^2(PY) < \lambda \widehat{\sigma}_\ell^2$ for all $\ell > k$. That is, $\widehat{k}$ has an even more convenient closed form

$$\widehat{k} = \max\left\{0 \le k \le K : \; d_k^2(PY) \ge \lambda \widehat{\sigma}_k^2 \; \right\} = \sum_{k=1}^{K} 1\left\{d_k^2(PY) \ge \lambda \widehat{\sigma}_k^2 \; \right\}. \tag{10}$$

Summarizing, we have shown the following result.

Theorem 3. *There exists a unique minimizer $\widehat{k}$ of (3), given by (10), such that $\widehat{\sigma}_k^2$ is monotone decreasing for $k \le \widehat{k}$, and monotone increasing for $k \ge \widehat{k}$.*

It is interesting to compare the choice $\widehat{k}$ in (10) with $\widehat{r}$ in [7]. In that paper, it is shown that (1) is equivalent with

$$\min_k \left\{\|Y - (PY)_k\|^2 + \mu k\right\} \tag{11}$$

based on the truncated singular value decomposition UD_kV^T of the projection $PY = UDV^T$ with $D_k = \text{diag}^\star(D_{11}, \dots, D_{kk}, 0, \dots, 0)$. Furthermore, [7] uses this formulation to derive a closed form for $\widehat{r}$, to wit,

$$\widehat{r} = \sum_{k \ge 1} 1\{d_k^2(PY) \ge \mu\} \tag{12}$$

based on the singular values $d_1(PY) \ge d_2(PY) \ge \cdots$ of the projection PY. The main difference between (10) and (12) is that $\widehat{k}$ counts the number of singular values of PY above a *variable* threshold, while $\widehat{r}$ counts the number of singular values of PY above a *fixed* threshold. Another difference is that the fixed threshold is proportional to the unknown variance

σ^2, while the variable threshold is proportional to $\widehat{\sigma}_k^2$, which can be thought of as an estimate of σ^2 only for k close to r.

To further illustrate the existence and uniqueness of $\widehat{k}$, we perform one experiment to show how the $d_k(PY)$ and $\widehat{\sigma}_k$ vary across different k, see Figure 2. The plot first displays the monotone property of $\widehat{\sigma}_k$ for $k \leq \widehat{k}$ and $k \geq \widehat{k}$. It also justifies the definition of $\widehat{k}$ in (10) since the rank which minimizes $\widehat{\sigma}_k$ is exactly what we defined.

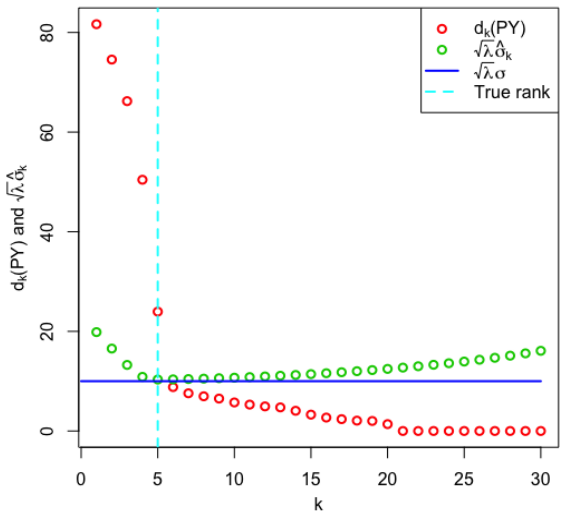

Figure 2: Plot of $d_k(PY)$ and $\sqrt{\lambda}\widehat{\sigma}_k$ vs k. In the experiment, we used $n = 150, m = 30, p = 20, q = 20, r = 5, \eta = 0.1, b_0 = 0.2$ and $\sigma^2 = 1$, using the notation and simulation setup of Section 5.3.

3 Rank consistency

3.1 The null case $XA = 0$

We treat the case $XA = 0$ separately as the case $XA \neq 0$ requires a lower bound on the non-zero singular values of XA.

Theorem 4. *Assume $XA = 0$. Then, on the event $\{d_1^2(PE) \leq \lambda\widehat{\sigma}^2\}$, we have $\widehat{k} = 0$.*

We particularize Theorem 4 to the case where the entries of E are independent $N(0, \sigma^2)$. In that case, general random matrix theory and Borel's inequality for suprema of Gaussian

processes, respectively, give $\mathbb{E}[d_1(PE)] \le \sigma(\sqrt{m}+\sqrt{q})$ and

$$\mathbb{P}\{d_1(PE) \ge \mathbb{E}[d_1(PE)] + \sigma t\} \quad \le \quad \exp(-t^2/2) \quad \text{for all } t > 0, \tag{13}$$

see Lemma 3 in Bunea, She and Wegkamp [7]. Moreover, $(nm)\widehat{\sigma}^2$ has a central χ^2_{mn} distribution, so that general tail bounds (Johnstone [21]) yield

$$\mathbb{P}\left\{\widehat{\sigma}^2 \le \sigma^2(1-\varepsilon)\right\} \quad \le \quad \exp\left(-mn\varepsilon^2/4\right), \quad 0 \le \varepsilon < 1 \tag{14}$$

$$\mathbb{P}\left\{\widehat{\sigma}^2 \ge \sigma^2(1+\varepsilon)\right\} \quad \le \quad \exp\left(-3mn\varepsilon^2/16\right), \ 0 \le \varepsilon < 1/2. \tag{15}$$

We immediately obtain the following corollary by using (13) – (15).

Corollary 5. *Assume that the entries of E are i.i.d. $N(0,\sigma^2)$. For any $\lambda > (\sqrt{m}+\sqrt{q})^2$, we have $\mathbb{P}\{\widehat{k}=0\} \to 0$ as $nm \to \infty$ and $m+q \to \infty$.*

The convergence rate in Corollary 5 is exponentially fast in nm and $m+q$.

3.2 The general case $XA \neq 0$

Theorem 6. *On the event,*

$$d_1^2(PE) \le \lambda\widehat{\sigma}_r^2 := \frac{\|Y-(PY)_r\|^2}{(nm/\lambda)-r} \tag{16}$$

we have $\widehat{k} \le r$.
In addition, for $1 \le s \le r$, on the event

$$d_s(XA) \ge d_1(PE) + \sqrt{\lambda}\widehat{\sigma}_r \tag{17}$$

intersected with the event (16), we further have $\widehat{k} \in [s,r]$.

While the quantity $\widehat{\sigma}_r^2$, used in the events (16) and (17) is a natural one in this problem, it depends on the unknown rank r. It turns out that quantifying $\widehat{\sigma}_r^2$ is not trivial. By properly choosing λ and by imposing a very mild rank constraint, however, we can give a more deterministic condition. Our rank constraint assumes that there exists $\delta_1, \delta_2 \in (0,1)$ such that the true rank r satisfies

$$0 \le r \le R := \begin{cases} K, & \text{if} \quad nm \ge (1+\delta_1)\lambda N \\ (1-\delta_2)K, & \text{otherwise} \end{cases} \tag{18}$$

The concern is to avoid a near zero denominator in criterion (3) at the true rank r. In particular, this suggests that for $nm/(nm - \lambda r) \leq \rho = \frac{1+\delta_1}{\delta_1} \wedge \frac{1}{\delta_2}$ with any $\delta_1 > 0$ and $\delta_2 \in (0, 1)$. The proof is shown in the appendix. For instance, taking $\delta_1 = 1/3$ and $\delta_2 = 1/4$ gives $\rho = 4$.

Corollary 7. *Under the rank constraint (18), on the event*

$$\{2d_1^2(PE) \leq \lambda\widehat{\sigma}^2\} \cap \{d_s(XA) \geq 2\widehat{\sigma}\sqrt{\rho\lambda}\},$$

for $1 \leq s \leq r$, we have $\widehat{k} \in [s, r]$.

We emphasize that the event $\{\lambda \geq 2d_1^2(PE)/\widehat{\sigma}^2\}$ itself guarantees $\hat{k} \leq r$. This is totally free of σ^2 and doesn't require any signal to noise condtion. We can further particularize to the Gaussian setting.

Theorem 8. *Assume the entries of E are i.i.d. $N(0, \sigma^2)$ and that r satisfies (18). Assume further that, for some constant $C > 2$, $d_s(XA) \geq C\sigma\sqrt{\rho\lambda}$ with $\lambda = C(\sqrt{m} + \sqrt{q})^2$. Then $\mathbb{P}\{s \leq \widehat{k} \leq r\} \to 1$ as $mn \to \infty$ and $m + q \to \infty$.*
If, in particular, $d_r(XA) \geq C\sigma\sqrt{\rho\lambda}$, then $\mathbb{P}\{\widehat{k} = r\} \to 1$ as $mn \to \infty$ and $m + q \to \infty$.

The convergence rate in Theorem 8 is exponentially fast in nm and $m + q$. This shows that the above procedure is highly accurate, which is confirmed in our simulation study. From the oracle inequality in the appendix, the prediction error $\|X\widehat{A}_{\widehat{k}} - XA\|^2$, for $s \leq \widehat{k} \leq r$, differs only within some constant levels of $d_1^2(PE)$ by which we denote the noise level.

4 Self-tuning procedure

From the main Theorem 6 in Section 3, we need to take λ, or in fact $\lambda\widehat{\sigma}_r^2$, as $\widehat{\sigma}_r^2$ depends on λ, large enough to prevent overfitting, that is, to avoid selecting a $\widehat{k}$ larger than the true rank r. On the other hand, we would like to keep λ small to be able to detect a small signal level. Indeed, (16) in Theorem 6, states that we need

$$d_1(PE) \leq \sqrt{\lambda}\widehat{\sigma}_r = \sqrt{\frac{\|Y - (PY)_r\|^2}{nm/\lambda - r}}. \tag{19}$$

The term on the right is decreasing in λ, so we should choose λ as small as possible such that $\sqrt{\lambda}\widehat{\sigma}_r$ is close to $d_1(PE)$. Without any prior knowledge on r, it is difficult to find the optimal

choice for λ. However, Corollary 7 tells us that an initial λ_0 satisfying $\{2d_1^2(PE) \leq \lambda_0\widehat{\sigma}^2\}$ yields an estimated rank $\widehat{k}_0$ with $\widehat{k}_0 \leq r$. Our idea is to use this lower bound $\widehat{k}_0$ for r to reduce our value λ_0 to λ_1. This, in turn, will yield a possibly larger estimated rank $\widehat{k}_1$, which still obeys $\widehat{k}_1 \leq r$. More precisely, we define

$$\widehat{\lambda}_0 := 2d_1^2(PE)/\widehat{\sigma}^2 \tag{20}$$

$$\widehat{k}_0 := \arg\min_k \frac{\|Y-(PY)_k\|^2}{nm-\widehat{\lambda}_0 k} \tag{21}$$

as starting values, and if $\widehat{k}_0 \geq 1$, for $t \geq 0$, we update

$$\widehat{\lambda}_{t+1} := \frac{nm}{\|E-(PE)_{(2\widehat{k}_t)\wedge N}\|^2/\widehat{U}_t + \widehat{k}_t} \tag{22}$$

$$\widehat{k}_{t+1} := \arg\min_k \frac{\|Y-(PY)_k\|^2}{nm-\widehat{\lambda}_{t+1} k} \tag{23}$$

with

$$\widehat{U}_t = d_1^2(PE) \vee \left\{ d_{2\widehat{k}_t-1}^2(PE) + d_{2\widehat{k}_t}^2(PE) \right\},$$

and the convention $d_j^2(PE) := 0$ for $j > N$. The procedure stops when $\widehat{k}_{t+1} = \widehat{k}_t$. We observe that (22) is free of σ^2 and the expected values of $\|E-(PE)_{(2\widehat{k}_t)\wedge N}\|^2$ and $\widehat{U}_t$ can be numerically evaluated by Monte Carlo simulations, provided the distribution of E is known. In particular, we will provide an analytical expression in the next section under the assumption that the entries E_{ij} are i.i.d. $N(0,\sigma^2)$, while various concentration inequalities guarantee that these bounds hold with overwhelming probability. All proofs are rather technical and deferred to the appendix.

NOTATION. In (22), we write $\widehat{\lambda}_{\widehat{k}_t}$ by $\widehat{\lambda}_{t+1}$ for simplicity and succinctness. We will stick to $\widehat{\lambda}_t$ when it does not cause confusion. In some cases, for clarification, we will write $\widehat{\lambda}_{\widehat{k}_t}$ explicitly instead.

Proposition 9. *We have $\widehat{\lambda}_t > \widehat{\lambda}_{t+1}$ and $\widehat{k}_t \leq \widehat{k}_{t+1}$ for $t = 0, 1, \ldots, T-1$. Moreover $\widehat{k}_t \leq r$ for all $t = 0, 1, \ldots, T$.*

This means that the sequence $\widehat{\lambda}_t$ is decreasing, implying a milder condition on the signal, while the sequence $\widehat{k}_t$ is increasing. These values are carefully chosen to maintain $\widehat{k}_t \leq r$. We refer to the proof for more explanations. The range of allowable ranks $\lfloor nm/(\widehat{\lambda}_t) \rfloor \wedge N$ is also

increasing in t. So as the iteration continues, the true rank could be selected even when it was out of the possible range at the beginning. This is verified in the later simulation.
In the following theorem, we write $\widehat{\lambda}_{k_t}$ explicitly to distinguish from $\widehat{\lambda}_{\widehat{k}_t}$ which is obtained based on $\widehat{k}_t$.

Theorem 10. *Assume $k_0 < k_1 < \cdots < k_T = r$ is a strictly increasing subsequence of $\{1, 2, \ldots, r\}$ of length $T + 1 \le r$. Define $\widehat{\lambda}_{k_{-1}} = \widehat{\lambda}_0$ as in (20) and $\widehat{\lambda}_{k_t}$ obtained from (22) by using k_t in lieu of $\widehat{k}_t$, for $t = 0, \ldots, T-1$. That is,*

$$\widehat{\lambda}_{k_t} \;:=\; \frac{nm}{\|E - (PE)_{(2k_t)\wedge N}\|^2/\widehat{U}_t + k_t} \tag{24}$$

Let $\widehat{k}_0 \le \widehat{k}_1 \le \cdots$ be the selected ranks from the procedure above, starting with the same $\widehat{\lambda}_0$ in (20). Then, on the event

$$d_{k_t}(XA) \ge d_1(PE) + \frac{\|Y - (PY)_r\|}{\sqrt{mn/\widehat{\lambda}_{k_{t-1}} - r}}, \qquad t = 0, \ldots, T, \tag{25}$$

there exists $s \le T$ such that $\widehat{k}_s = r$.

Another way to interpret Theorem 10 is to think $\{k_t\}_{t=0}^T$ as the sequence to guarantee rank recovery. The ideal case is $\{k_t\}_{t=0}^T = \{r\}$ which leads to a one-step recovery, but requires a comparatively stronger signal to noise condition in (25). At the other extreme, it could take r steps to recover the true rank. We emphasize that the latter case requires the *mildest* signal to noise condition (25) on the singular values of XA to be greater than a decreasing sequence of thresholds. In particular, we observe that the condition

$$d_r(XA) \ge d_1(PE) + \frac{\|Y - (PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{k_{T-1}} - r}}$$

at the true rank lowers the previous one in (17) which is the same formula with $\widehat{\lambda}_{k_{T-1}}$ replaced by the larger value $\widehat{\lambda}_0$.

4.1 Self-tuning procedure in case of Gaussian entries

If E has i.i.d. $N(0, \sigma^2)$ entries, a deterministic bound for updating λ is derived as follows. We define $M := m \vee q$ corresponding to $N = q \wedge m$. For given $\widetilde{k}_t \ge 1$ and $0 < \varepsilon < 1$, we set

$$\widetilde{\lambda}_{t+1} = \begin{cases} \Big(nm\widetilde{U}_t\Big)\Big/\Big((1-\varepsilon)nm - \widetilde{R}_t + \widetilde{k}_t\widetilde{U}_t\Big) & \text{if } 2\widetilde{k}_t < N; \\ \Big\{nm(\sqrt{q}+\sqrt{m})^2\Big\}\Big/\Big\{(1-\varepsilon)(n-q)m + \widetilde{k}_t(\sqrt{q}+\sqrt{m})^2\Big\} & \text{otherwise} \end{cases} \tag{26}$$

with

$$\widetilde{U}_t := \max\left\{(\sqrt{m}+\sqrt{q})^2, \left(\sqrt{M}+\sqrt{N-2\widetilde{k}_t+2}\right)^2+\left(\sqrt{M}+\sqrt{N-2\widetilde{k}_t+1}\right)^2\right\} \tag{27}$$

and

$$\widetilde{R}_t := \min\left\{\sum_{j=1}^{2\widetilde{k}_t}\left(\sqrt{M}+\sqrt{N-j+1}\right)^2 + r, qm - \sum_{j=1}^{2\widetilde{k}_t}\left(\sqrt{M}-\sqrt{j}\right)^2\right\}. \tag{28}$$

Proposition 11. *For any given $\{k_t\}_{t\geq 0}$ with $1 \leq k_t \leq r$, we let $\widehat{\lambda}_{t+1}$ and $\widetilde{\lambda}_{t+1}$ obtained from (22) and (26) by using (the same) k_t in place of $\widehat{k}_t$ and $\widetilde{k}_t$, respectively. In addition, we let $\widetilde{\lambda}_0 = 2(\sqrt{m}+\sqrt{q})^2$ and $\widehat{\lambda}_0 = 2d_1^2(PE)/\widehat{\sigma}^2$, then we have $\mathbb{P}\{\widetilde{\lambda}_t \geq \widehat{\lambda}_t\} \to 1$, for all $t \geq 0$.*

Proposition 11 indicates $\widetilde{\lambda}_t$ is a deterministic (upper) bound for $\widehat{\lambda}_t$ with overwhelming probability for given $\widetilde{k}_t = \widehat{k}_t$. The following proposition states that $\widetilde{\lambda}_t$ itself is still decreasing and, meanwhile, large enough to prevent overfitting.

Proposition 12. *Let $\widetilde{k}_t$ be the selected rank from (23) by using $\widetilde{\lambda}_t$ chosen according to (26) for $t \geq 0$ starting with $\widetilde{\lambda}_0 = 2(\sqrt{m}+\sqrt{q})^2$. Then we have $\widetilde{\lambda}_{t+1} \leq \widetilde{\lambda}_t$ and $\mathbb{P}\{\widetilde{k}_t \leq r\} \to 1$ for all t.*

The convergence rate in Proposition 12 is exponentially fast in nm and $m+q$. Armed with this, we propose the following Algorithm 1 for Self-tuning Rank Selection (SRS) under Gaussianity.

Algorithm 1 Self-tuning Rank Selection

procedure
  initialize $\widetilde{\lambda}_0 = 2(\sqrt{m}+\sqrt{q})^2$
  $\widetilde{k}_0 \leftarrow$ selected from (21) by using $\widetilde{\lambda}_0$
  if $\widetilde{k}_0 = 0$ **then**
    return $\widetilde{k}_0$
  else
    repeat
      $\widetilde{\lambda}_{i+1} \leftarrow$ obtained from (26) by using $\widetilde{k}_i$
      $\widetilde{k}_{i+1} \leftarrow$ selected from (23) by using $\widetilde{\lambda}_{i+1}$ over $\widetilde{k}_i, \widetilde{k}_i+1, \cdots, K_{\widetilde{\lambda}_{i+1}}$
    until $\widetilde{k}_{i+1} = \widetilde{k}_i$
    return $\widetilde{k}_i$

Propositions 11 and 12 ensure the decreasing sequence of $\widetilde{\lambda}_i$ which leads to a non-decreasing sequence of updated rank $\widetilde{k}_i$. At the same time, it also guarantees that the newly selected ranks are no greater than r. Therefore, if a proper signal to noise condition is met, Algorithm 1 guarantees recovery of the true rank.

Regarding the computational complexity, we point out that SRS has almost the same level of computational complexity as those methods in (1) and (3) without self-tuning. This is due to that the computationally expensive singular value decomposition only needs to be computed once. In addition, finding the new rank based on updated λ only needs to start from the previous selected rank. This avoids a lot of extra computation.

In addition to using the deterministic bounds to update λ, we can use Monte Carlo simulations to numerically evaluate the expectation of the random part in (22). For settings with smaller dimensions, the performance of analytical bounds are not as good as that for large dimensions. As Giraud [17] pointed out, the Marchenko-Pastur distribution [24] could also be used as an accurate approximation in high dimensional settings.

If we further define $\widetilde{\rho}_i := (nm)/(nm - \widetilde{\lambda}_i r)$ for $i = 0, 1, \ldots$, we can derive a more explicit signal to noise condition in the following theorem.

Theorem 13. *Assume E has i.i.d. $N(0, \sigma^2)$ entries. For any subsequence $k_0 < \cdots < k_T = r$ of $\{1, 2, \ldots, r\}$ with $T \leq r - 1$, we let $\widetilde{\lambda}_{k_{-1}} = \widetilde{\lambda}_0 = 2(\sqrt{m} + \sqrt{q})^2$ and denote by $\widetilde{\lambda}_{k_t}$ the updated λ according to (26) by using k_t, for $t = 0, 1, \ldots, T - 1$. Suppose $\widetilde{k}_0, \widetilde{k}_1, \ldots$ are the selected ranks from Algorithm 1. Then on the event,*

$$d_{k_t}(XA) \geq C\sigma\sqrt{\widetilde{\rho}_t \widetilde{\lambda}_{k_{t-1}}}, \qquad t = 0, 1, \ldots, T. \tag{29}$$

for some $C > 2$, there exists $s \leq T$ such that $\mathbb{P}\{\widetilde{k}_s = r\} \to 1$.

The convergence in Theorem 13 is exponentially fast in nm and $m+q$. The sequence $\{\widetilde{k}_t\}_{t=0}^{s}$ here has the same role as $\{\widehat{k}_t\}_{t=0}^{s}$ in Theorem 10. The length of this sequence is the number of steps we need to recover r, which depends on the signal to noise condition in (29).
The signal to noise condition (29) can be particularized at r as

$$d_r(XA) \geq C\sigma\sqrt{\frac{nm}{nm/\widetilde{\lambda}_{k_{T-1}} - r}}$$

which lowers the one in Theorem 8 provided $\widetilde{\lambda}_{k_{T-1}} \leq \widetilde{\lambda}_{k_{-1}}$.

5 Empirical study

Through Sections 6.1 – 6.5, we perform simulations to compare the methods discussed in this paper with other methods. In Section 6.6, we perform an additional simulation to check the tightness of signal to noise condition in (17).

5.1 Methods and notations

We first introduce different methods which appreared in our simulation. Bunea, She and Wegkamp [7] proposed the method in (1) to select the optimal rank by using $\mu = Cd_1^2(P\widetilde{E})\widetilde{\sigma}^2$, where $\widetilde{E}$ has i.i.d. $N(0,1)$ entries and $\widetilde{\sigma}^2$ in (2) is the unbiased estimator of σ^2. $C > 1$ is the leading constant to be specified. A deterministic upper bound which could be used instead is $C(\sqrt{m}+\sqrt{q})^2\widetilde{\sigma}^2$. The second method is proposed by Giraud [17] to minimize

$$\frac{\|Y-(PY)_k\|^2}{nm-1-C\left(\mathbb{E}[\|G\|_{(2,r)}]\right)^2} \tag{30}$$

with $\|G\|_{2,r}^2 = \sum_{k=1}^r d_k^2(G)$ and some tuning parameter $C > 1$. It does not require to estimate σ^2 but still needs the selection of leading constant C. Giraud [17] recommended to use $C = 2$ based on Birgé and Massart [5]. The method proposed and studied in the current paper is also free of σ^2. The one without self-tuning also requires selecting $C > 2$ for $\lambda = Cd_1^2(PE)/\widehat{\sigma}^2$ which could be upper bounded by $C(\sqrt{m}+\sqrt{q})^2$. However, there is no reason for one choice of the tuning parameter being globally optimal. We verified that this is indeed the case by simulations. Another option for choosing the tuning parameter is to use k-fold Cross-Validation (CV). However, there is no theoretical guarantee of the feasibility for Cross-Validation, especially if the rows $X_{i\cdot}$ of X are non-i.i.d. In contrast, our proposed self-tuning method is completely devoid of choosing tuning parameter and estimating σ^2.

Notation. We use BSW to denote the method proposed in Bunea, She and Wegkamp [7], KF to denote the method proposed in Giraud [17], GRS and SRS to denote the method without and with self-tuning, respectively, proposed in this paper. We further use BSW-K to denote BSW with specified leading constant K and in the same way, we define KF-2

and GRS-2. Finally, we denote by KF-CV the KF method with leading constant chosen by Cross-Validation.

5.2 Task description

We divide the whole simulation into four parts. In the first part, we focus on the comparison of BSW-K by using different leading constant K. We show that there is no optimal leading constant which works for all r. In the second part, we show the improvement of SRS over GRS in terms of requiring a smaller signal to noise ratio (SNR) and enlarging the range of possible selected ranks. The third part compares the performance of SRS and BSW-K in three dimension settings. The last part presents an example where KF fails to recover the true rank.

5.3 Simulation setup

In general, the simulation setup lies in three different categories. The first one is the more favorable case where the sample size n is comparatively larger than the number of covariates p. The other two are high-dimensional in which n is smaller than p, hence more challenging. The second one of these high-dimensional cases focus on the worst case when n is close to q for a moderate m.

Our experiments are inspired by those of Bunea, She and Wegkamp [7]. The $n \times p$ design matrix X is generated by independently drawing n times p-dimensional Gaussian vectors with mean zero and covariance matrix specified by $\Sigma_{i,j} = \eta^{|i-j|}$ with $\eta \in (0,1)$, for $i, j = 1, \ldots, p$. We always set $\sigma^2 = 1$. The regression coefficient matrix A is given by $A = b_0 M_{p\times r} M_{r\times m}$ where the entries of M are i.i.d. $N(0,1)$. As before, r denotes the true rank of A satisfying $r \leq N = q \wedge m$.

The difference with [7] lies in the way of varying the Signal-to-Noise-Ratio (SNR) which is denoted by $d_r(XA)/d_1(PE)$. Instead of using the combinations of different η and b_0, we vary the SNR by generating A over different true ranks. Specifically, we first generate the fixed design matrix X based on given η. Then for each r in the range of true ranks, we generate A according to the previous scheme for given b_0. Finally, we repeat 200 times for each pair of (X, A) by generating Gaussian matrices and using different methods to select

the optimal rank. Besides recording the rank recovery rate and the mean selected rank in the 200 replications, we calculate the SNR for each pair of (X, A).

Monte-Carlo simulations can be used for all of those methods to numerically evaluate the random penalty terms. We found, from extensive simulations, that Monte-Carlo simulations are to be preferred over deterministic (upper) bounds in general, especially in low dimensional settings. Therefore, throughout our simulation, we use Monte-Carlo simulations for updating λ for SRS, and to keep consistency, for approximating $d_1^2(P\widetilde{E})$ in the penalty term of the BSW method.

5.4 Experiment 1

In this part, we compare the rank recovery of BSW-K by varying K in $\{0.7, 0.9, 1.1, 1.3, 1.5\}$ with SRS as the benchmark in both low-dimensional and high-dimensional settings. In the low-dimensional setting, we consider the situation when $n = 150$, $m = 30$ and $p = q = 20$. We vary the true rank from 0 to 20 and set $\eta = 0.1$. For b_0, we choose from $\{0.15, 0.20, 0.25\}$ to illustrate, more clearly, the effect of rank r on recovery rate. On the other hand, the high-dimensional setting uses $n = 100$, $m = 30$, $p = 150$, $q = 20$ and $\eta = 0.1$. The true rank r is changed from 0 to 20. We choose b_0 from $\{0.03, 0.05, 0.07\}$. The results of low-dimension and high-dimension are shown in Figures 3 and 4, respectively.

Result: Both figures clearly demonstrate that a smaller K, say 0.7 or 0.9, performs better when the true rank r is big, in the sense of requiring a smaller SNR, but tends to overfit for small r. In contrast, a larger K does a better job in preventing overfitting for small r but requires a larger SNR. It is interesting to see that the performance of BSW-K does depend on r. Note that by varying b_0, we separate the effect of SNR away from this phenomenon. Therefore, this strongly suggests that there is no optimal leading constant K for BSW to guarantee the performance for all possible r. On the other hand, SRS performs globally better and more stable than RSC-K in both settings. For smaller ranks, it prevents overfitting and meanwhile requires a smaller SNR than BSW-1.3 and BSW-1.5 which also prevent overfitting. In addition, it is also worth noting that the role of SNR for the rank recovery is striking. If SNR is too small, say smaller than 0.8, then the rank recovery completely fails. This justifies the signal to noise condition in (17) and is explained by the exponential tail bounds in our main results.

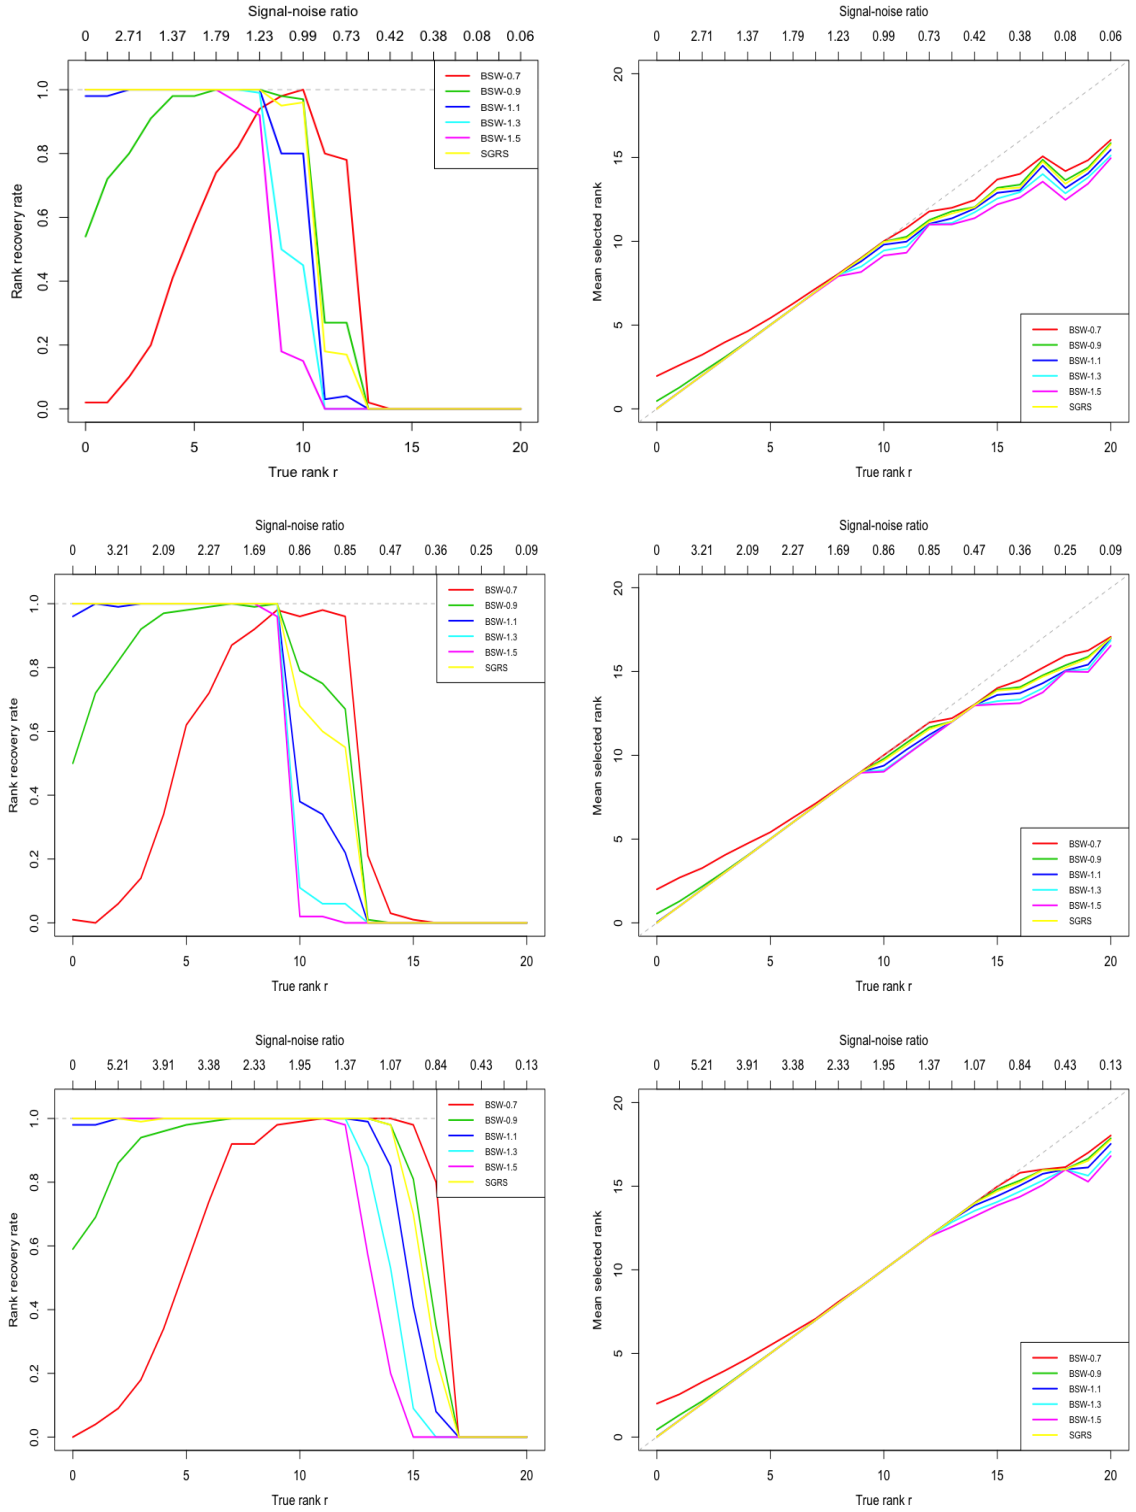

Figure 3: Performance of BSW-K and SRS in the low-dimensional setting of Experiment 1. b_0 is $0.15, 0.20$ and 0.25 for each row, respectively. The left one of each row is the plot of rank recovery vs r and the right is mean selected rank vs r.

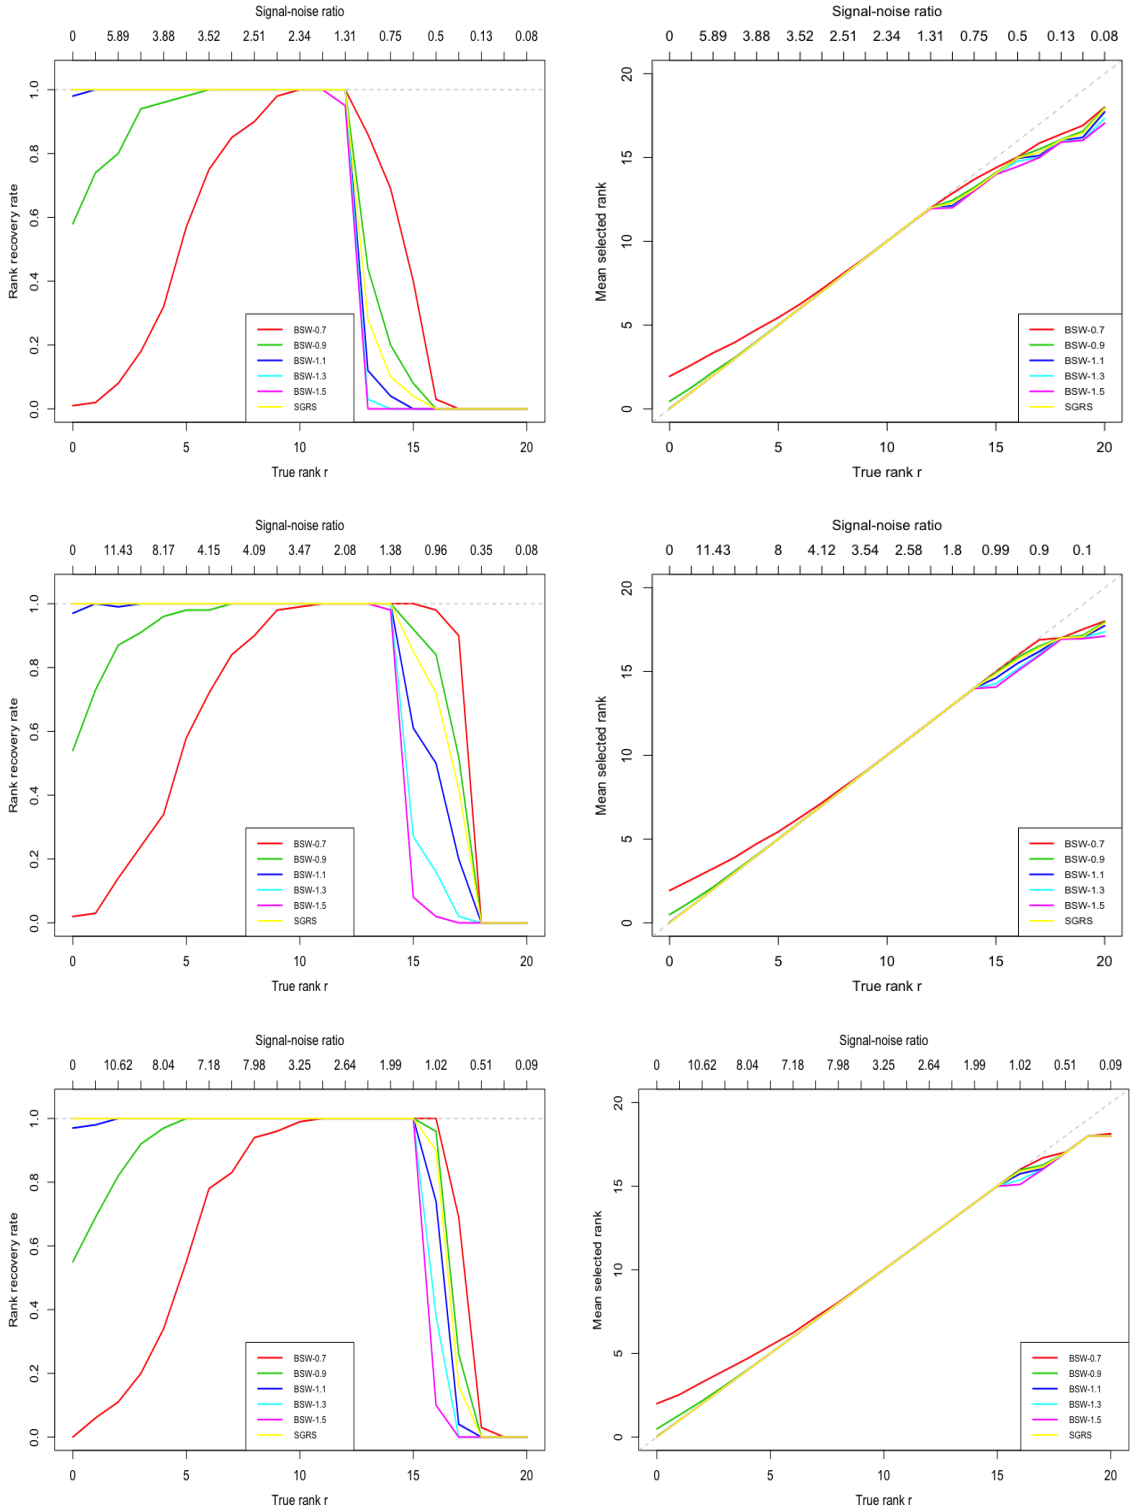

Figure 4: Performance of BSW-K and SRS in the high-dimensional setting of Experiment 1. b_0 is $0.03, 0.05$ and 0.07 for each row, respectively. The left one of each row is the plot of rank recovery vs r and the right is mean selected rank vs r.

5.5 Experiment 2

Here we compare the performance of GRS and SRS in the following three settings. The first one has the same low-dimensional setting as that in Experiment 1 with $b_0 = 0.25$ while the second one is using the same high-dimensional setting in Experiment 1 with $b_0 = 0.07$. The third setting focuses on the case when $nm \leq \lambda_0 N$ which incurs the rank constraint in (18). Specifically, we set $n = 50$, $m = 50$, $p = 300$, $q = 30$, $\rho = 0.1$ and $b_0 = 2$. The true rank r is ranging from 0 to 30. Note in this setup, $K_{\lambda_0} = 7$.

RESULT: The plots of the first two settings in Figure 5 indicate that the required SNR of SRS is reduced to around 1 from that of GRS which is about 2. The plot in Figure 6 shows that the largest recoverable rank K_λ is extended by SRS through reducing λ. GRS fails to recover $r > K_{\lambda_0}$ whereas SRS perfectly recovers all possible ranks. This verifies that when n is not too small, SRS can get rid of the rank constraint (18) as λ becomes increasingly smaller. It is worth mentioning that λ reduces from 198 to 83 and 66 for the first two cases, respectively, and from 315 to 69 for the third one.

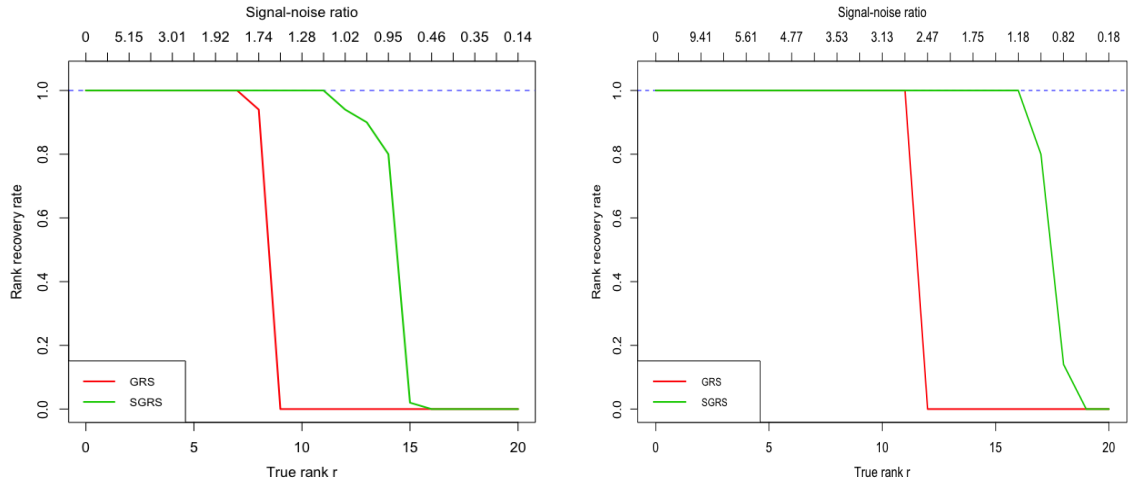

Figure 5: Performance of rank recovery rate vs true rank r for GSR and SRS in the first two dimensional settings of Experiment 2.

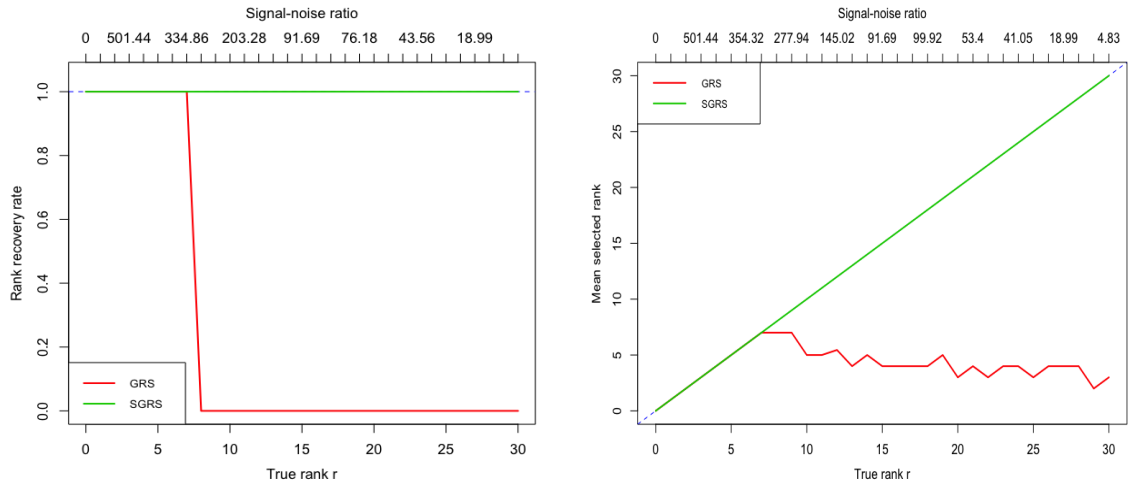

Figure 6: Performance of rank recovery rate vs true rank r for GSR and SRS in the third dimensional setting of Experiment 2.

5.6 Experiment 3

Based on the previous results, we choose BSW-1.1 and BSW-1.3 as the representative methods for BSW-K since they seem to be a better compromise. No KF method is considered here since later in Experiment 4 we will present an example to show that it does not guarantee rank consistency. Since SRS always outperforms GSR, we neither involve GRS. Figures 3 and 4 have already shown the advantage of SRS over BSW-K for both the low-dimensional case and the high-dimensional when n is not too small. Therefore, in this part, we only focus on the worst case, that is, when n is close to q. Besides, we also verify the feasibility of SRS when $n = q$ and BSW is infeasible, by setting different dimensions.

For the comparison between BSW-K and SRS, we focus on four cases where we set the same $n = 150, m = 30, p = 200, \eta = 0.1$ and $b_0 = 0.011$, but vary q in $\{143, 145, 147, 149\}$. The true rank r is from 0 to 22. The plot of recovery rate in this setting is shown in Figure 7.

Result: It is clear to see that, given a moderate m, the closer n to q is, the worse BSW-K performs. This is as expected since the estimation of σ^2 becomes problematic when $(n-q)m$ is not large. The same problem for choosing different K is still noticeable. Specifically, BSW-1.1 requires a smaller SNR however overfits more than BSW-1.3 at small r. On the other

hand, SRS performs perfectly as long as r lies in the allowable range. The upper recoverable rank for SRS in this case is 11 to 13 depending on the particular choice of q. Note that in this case, $r \leq q \wedge m = 30$.

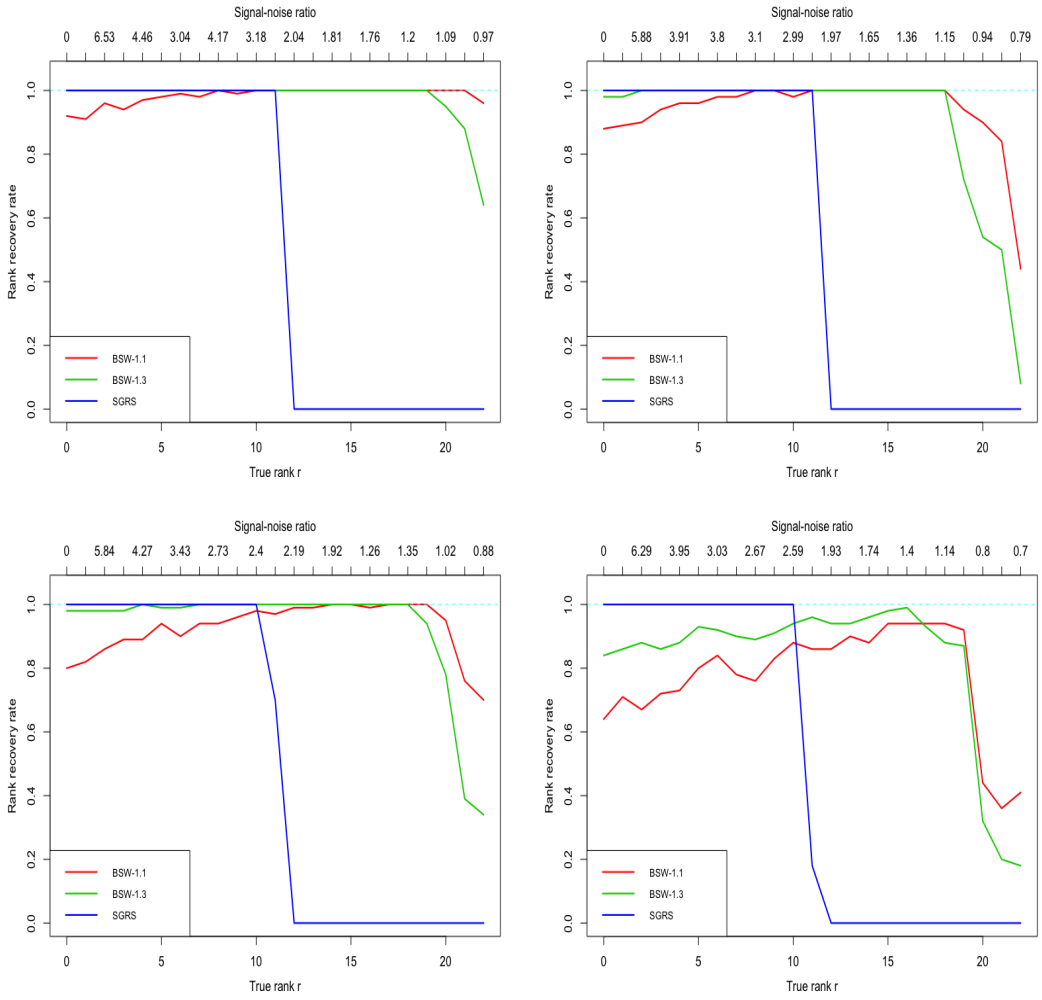

Figure 7: Plots of rank recovery rate for BSW-1.1, BSW-1.3 and SRS in Experiment 3. The value of q varies in $\{143, 145, 147, 149\}$ from left to right and top to bottom, correspondingly.

For the case when $n = q$, we test the recovery rate of SRS based on nine settings with the same $p = 200$, $b_0 = 0.011$ and $\eta = 0.1$, but varying n, q and m, correspondingly, in $\{50, 100, 150\}$ and $\{50, 125, 200\}$. We also vary r from 0 to 50. In all of these settings, we set

SNR large enough, say above 3, to eliminate the effect of SNR on rank recovery. Figure 8 displays the performance of SRS for all combinations of $n(q)$ and m. SRS performs perfectly as long as r is within the recoverable range. It is also worth noticing that when $n = q$, the allowable range also depends on m. The larger m is, the larger the range is. So when $n = q$, as long as r is not very large or m is not quite small comparing to q, SRS has a guaranteed performance.

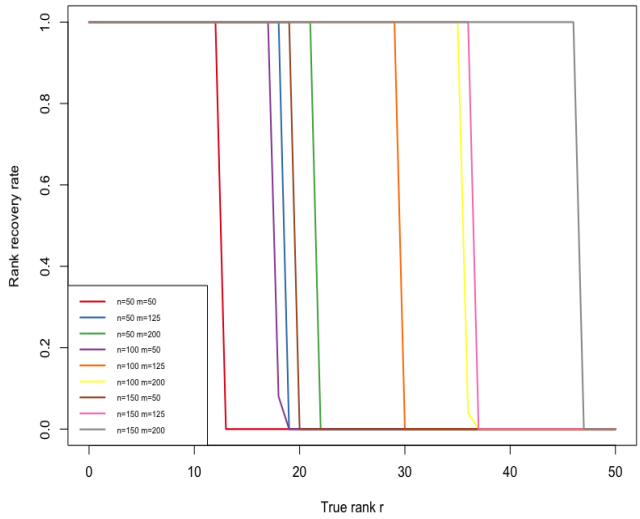

Figure 8: Rank recovery rate for SRS when $n = q$ in Experiment 3.

5.7 Experiment 4

This setting is slightly different from the low dimensional one in Experiment 1 and aims to give an example of an easy situation where KF fails to select the correct rank. Here we have $n = 300$, $m = 40$ and $p = q = 35$. For illustration, we only vary the true rank from 10 to 35. We set $\eta = 0.1$ and $b_0 = 20$. Note $b_0 = 20$ ensures a large SNR which should be the most ideal case for rank recovery. We compare the performance of KF-2, KF-CV and SRS. The plot of rank recovery and mean selected rank versus the true rank for different methods are shown in Figure 9. From the result, it is clear that neither KF-2 nor KF-CV consistently recovers the true rank while SRS perfectly recovers the true rank. But since KF was not developed for rank recovery, nor did it claim to have this property, this does not contradict the results in Giraud [17].

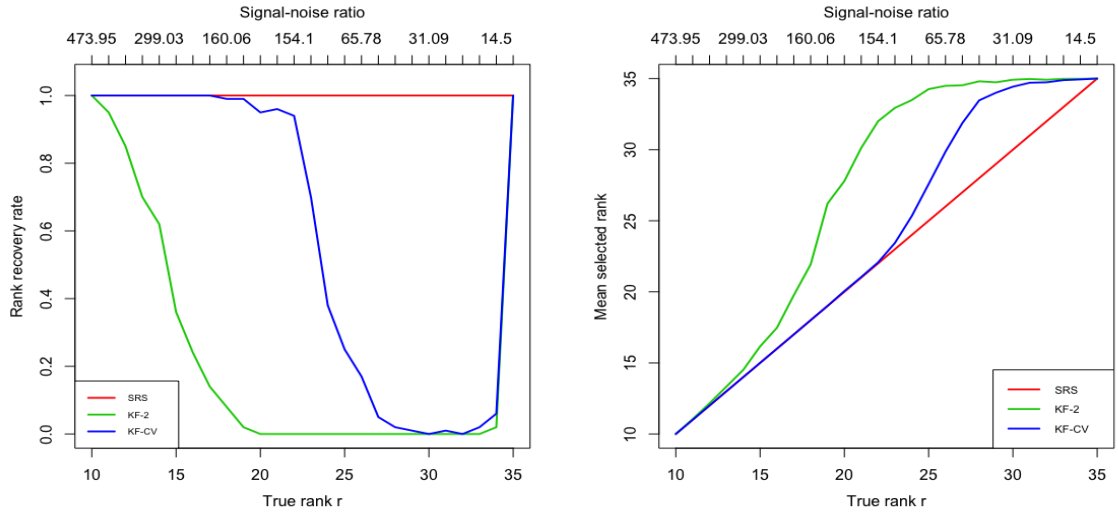

Figure 9: Plots of performance for KF methods and SRS in Experiment 4. The left is the plot of rank recovery vs r and the right is mean selected rank vs r.

Conclusion

We can draw the following conclusions:

- In general, SRS outperforms than BSW-K in both low-dimensional and high-dimensional settings in both the rank recovery rate and stability. The performance of BSW-K by using different leading constant does depend on the true rank r. This leads choosing the tuning parameter problematic as we have shown there is no globally optimal K for BSW-K. On the other hand, SRS is globally stable for all settings as long as r lies in its allowable range.

- In the most challenging settings of Experiment 4, the advantage of SRS over BSW-1.1 and BSW-1.3 is enlarged when r lies within the allowable range. This is as expected since the estimation of σ^2 is problematic.

- When $n = q$, SRS still maintains its stable consistency. The only drawback of SRS is the rank constraint when nm is close to $m + q$. But this condition only depends on q and m rather than p. Therefore, SRS only fails when r is large and both of n and m is small comparing to q, which is a rare case. Clearly, reduced rank regression makes only sense for relatively small r.

5.8 Tightness check of signal to noise condition

Akin to the discussion in [7, Section 4.2, p. 1303], we can empirically verify the tightness of signal to noise condition in (17). Specifically, from (10), we have

$$\{\widehat{k} \neq r\} = \{d_{r+1}(PY) \geq \sqrt{\lambda}\widehat{\sigma}_r\} \cup \{d_r(PY) \leq \sqrt{\lambda}\widehat{\sigma}_r\}. \tag{31}$$

By using identity (31) and Weyl's inequality, it is not hard to obtain

$$\mathbb{P}\{\widehat{k} \neq r\} \geq \mathbb{P}\left\{d_r(XA) + d_1(PE) \leq \sqrt{\lambda}\widehat{\sigma}_r\right\}.$$

Hence we conclude that $\mathbb{P}\{d_r(XA) \leq -d_1(PE) + \sqrt{\lambda}\widehat{\sigma}_r\} > 0$ implies $\mathbb{P}\{\widehat{k} = r\} < 1$. Again, using identity (31) and Weyl's inequality, we observe that

$$\mathbb{P}\{\widehat{k} \neq r\} \leq \mathbb{P}\left\{d_1(PE) \geq \min\left(\sqrt{\lambda}\widehat{\sigma}_r, d_r(XA) - \sqrt{\lambda}\widehat{\sigma}_r\right)\right\}. \tag{32}$$

Since $\mathbb{P}\{d_1(PE) \geq \sqrt{\lambda}\widehat{\sigma}_r\} \to 0$ for appropriately chosen λ, we claim that $d_r(XA)$ cannot deviate from $\sqrt{\lambda}\widehat{\sigma}_r$ much more than $d_1(PE)$. To empirically verify this conjecture, we generate different pairs of (X, A) through changing b_0, η, n, m, p, q and r. For each pair of (X, A), on the one side, we record the rth largest singular value of XA as $d_r(XA)$; on the other hand, we search along the grid of λ to find the largest λ such that minimizing (3) recovers the true rank, given the fact that $\sqrt{\lambda}\widehat{\sigma}_r$ is increasing in λ. Finally, we make the plots of $\lambda\widehat{\sigma}_r$ and $d_1(PE) \pm \sqrt{\lambda}\widehat{\sigma}_r$ versus $d_r(XA)$ for all pairs of (X, A) which are shown in Figure 10. This strongly confirms the tightness of the signal to noise condition in (17).

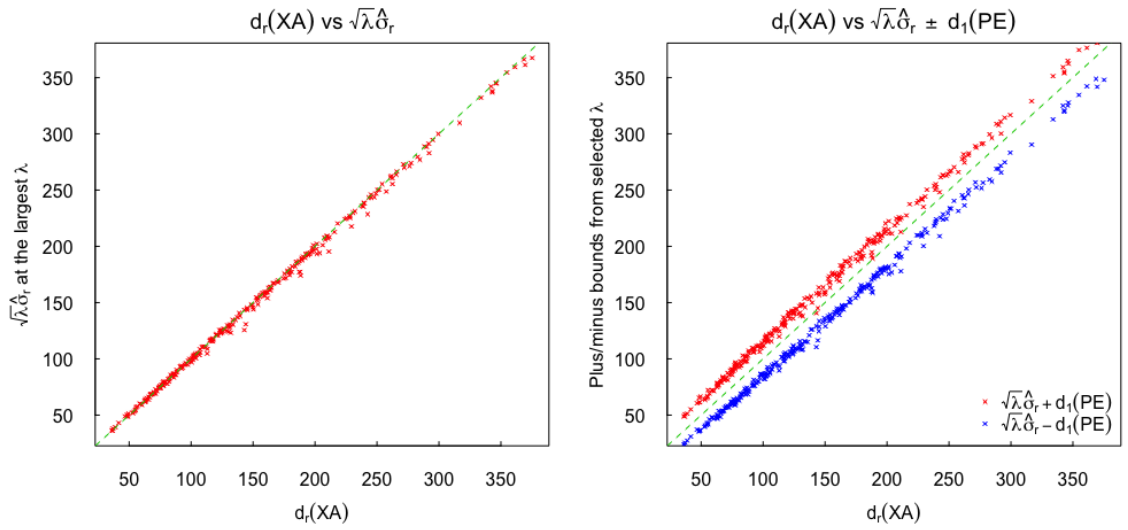

Figure 10: The plots of $\sqrt{\lambda}\widehat{\sigma}_r$ and $d_1(PE) \pm \sqrt{\lambda}\widehat{\sigma}_r$ versus $d_r(XA)$ for each pair of (X, A). λ is the largest one searched along the given grid which leads to the true rank recovery. The green dotted line has zero intercept and slope equal to 1.

Appendix A Proofs of Sections 2 and 3

A.1 Poof of Proposition 1

We first note that, since $(PY)_i = P(PY)_i$, see Giraud [18] (page 124), and Pythagoras' identity

$$\begin{aligned} \|Y - (PY)_i\|^2 &= \|Y - PY\|^2 + \|PY - (PY)_i\|^2 \\ &= \|Y - PY\|^2 + \sum_{k=i+1}^{j} d_k^2(PY) + \sum_{k>j} d_k^2(PY) \\ &= \|Y - (PY)_j\|^2 + \sum_{k=i+1}^{j} d_k^2(PY). \end{aligned} \tag{33}$$

Consequently, with $i < j$,

$$\begin{aligned} \widehat{\sigma}_i^2 \geq \widehat{\sigma}_j^2 &\iff \frac{\|Y - (PY)_i\|^2}{nm - \lambda i} \geq \frac{\|Y - (PY)_j\|^2}{nm - \lambda j} \\ &\iff \frac{\|Y - (PY)_j + \sum_{k=i+1}^{j} d_k^2(PY)}{(nm - \lambda j) + \lambda(j - i)} \geq \frac{\|Y - (PY)_j\|^2}{nm - \lambda j} \\ &\iff \frac{\sum_{k=i+1}^{j} d_k^2(PY)}{\lambda(j - i)} \geq \frac{\|Y - (PY)_j\|^2}{nm - \lambda j} = \widehat{\sigma}_j^2 \end{aligned}$$

using the simple fact that $(a+b)/(c+d) \geq a/c \iff (b/d) \geq (a/c)$ for any positive numbers a, b, c, d. This proves (4). Claim (5) follows from (4) by taking $i = j-1$. Finally,

$$\frac{d_j^2(PY)}{\lambda} \leq \frac{\|Y-(PY)_j\|^2}{nm-\lambda j} = \frac{\|Y-(PY)_{j-1}\|^2 + d_j^2(PY)}{nm-\lambda(j-1)+\lambda}$$

is equivalent with

$$\frac{d_j^2(PY)}{\lambda} \leq \frac{\|Y-(PY)_{j-1}\|^2}{nm-\lambda(j-1)} = \widehat{\sigma}_{j-1}^2$$

using the above elementary manipulation again and (6) follows. This completes our proof. □

A.2 Proof of Proposition 2

We first show (7). Suppose $\widehat{\sigma}_k^2 \leq \widehat{\sigma}_{k-1}^2$. We observe

$$\begin{aligned} \frac{1}{k-\ell}\sum_{j=\ell+1}^{k} d_j^2(PY) &\geq d_k^2(PY) && \text{by } d_1(PY) \geq d_2(PY) \geq \cdots \\ &\geq \lambda\widehat{\sigma}_{k-1}^2 && \text{by (5)} \end{aligned}$$

so that (4) implies $\widehat{\sigma}_k^2 \leq \widehat{\sigma}_\ell^2$ for all $\ell \leq k-1$. This proves the non-trivial direction of (7). Next, we show (8). Suppose $\widehat{\sigma}_k^2 \geq \widehat{\sigma}_{k-1}^2$. Then, following the same reasoning as above, we get

$$\frac{1}{\ell-k+1}\sum_{j=k-1}^{\ell} d_j^2(PY) \leq d_{k-1}^2(PY) < \lambda\widehat{\sigma}_{k-1}^2,$$

and $\widehat{\sigma}_\ell^2 \geq \widehat{\sigma}_{k-1}^2$ by (4). This completes our proof. □

A.3 Proof of Theorem 4

By Theorem 3, it suffices to show $d_1^2(PY) \leq \lambda\widehat{\sigma}_1^2$. This is equivalent to $d_1^2(PY) \leq \lambda\widehat{\sigma}_0^2$ by criterion (6) in Proposition 1. The latter, in turn, is equivalent to $d_1^2(PE) \leq \lambda\widehat{\sigma}^2$ as $XA = 0$ implies $d_1^2(PE) = d_1^2(PY)$ and $\widehat{\sigma}_0^2 = \widehat{\sigma}^2$. □

A.4 Proof of Corollary 5

We can write $\lambda := C(\sqrt{m}+\sqrt{q})^2$ for some $C = (1+C_0)^2/(1-C_1) > 2$ with $C_0 > 0$ and $0 < C_1 < 1$. By Theorem 4, (13) and (14), we have

$$\begin{aligned}
\mathbb{P}\{\widehat{k} \neq 0\} &= \mathbb{P}\left\{ d_1^2(PE) \geq \lambda \frac{\|E\|^2}{nm} \right\} \\
&\leq \mathbb{P}\left\{ d_1^2(PE) \geq (1+C_0)^2(\sqrt{m}+\sqrt{q})^2\sigma^2 \right\} \\
&\quad + \mathbb{P}\left\{ \|E\|^2 \leq (1-C_1)nm\sigma^2 \right\} \\
&\leq \exp\left\{ -C_0^2(\sqrt{m}+\sqrt{q})^2/2 \right\} + \exp\left\{ -C_1^2 nm/4 \right\},
\end{aligned}$$

which proves the claim. □

A.5 Proof of Theorem 6

By Proposition 1, for $\widehat{k} \leq r$, we need to show that

$$d_{r+1}^2(PY) \quad < \quad \lambda \widehat{\sigma}_{r+1}^2. \tag{34}$$

We observe that

$$d_{r+1}^2(PY) < \lambda \widehat{\sigma}_{r+1}^2 \quad \Longleftrightarrow \quad \lambda > d_{r+1}^2(PY)/\widehat{\sigma}_r^2$$

by statement (6). Again, an application of Weyl's inequality and observing that $d_{r+1}(XA) = 0$, we have, on the event (16),

$$\lambda \geq \frac{d_1^2(PE)}{\widehat{\sigma}_r^2} \geq \frac{d_{r+1}^2(PY)}{\widehat{\sigma}_r^2}$$

which is exactly what needed to be shown for (34).
To show $\widehat{k} \geq s$, it suffices to prove $d_s^2(PY) \geq \lambda \widehat{\sigma}_s^2$. Identity (33) gives

$$\begin{aligned}
\|Y - (PY)_s\|^2 &= \|Y - (PY)_r\|^2 + \sum_{j=s+1}^{r} d_j^2(PY) \\
&\leq \|Y - (PY)_r\|^2 + (r-s) d_{s+1}^2(PY).
\end{aligned}$$

Consequently,

$$d_s^2(PY) \quad \geq \quad \lambda \frac{\|Y - (PY)_s\|^2}{nm - \lambda s}$$

is implied by

$$d_s^2(PY) \quad \geq \quad \lambda \frac{\|Y - (PY)_r\|^2 + (r-s)d_s^2(PY)}{nm - \lambda s}$$

which in turn, after a little algebra, is seen to be equivalent to

$$d_s^2(PY) \quad \geq \quad \lambda \frac{\|Y - (PY)_r\|^2}{nm - \lambda r} = \lambda \widehat{\sigma}_r^2.$$

By Weyl's inequality, on the event (17), we have

$$d_s(PY) \geq d_s(XA) - d_1(PE) \geq \sqrt{\lambda}\widehat{\sigma}_r,$$

which shows $d_s^2(PY) \geq \lambda \widehat{\sigma}_s^2$. □

A.6 Proof of Corollary 7

Note the events in this corollary combined with

$$\widehat{\sigma}^2 \leq \widehat{\sigma}_r^2 \leq \rho \widehat{\sigma}^2 \tag{35}$$

imply (16) and (17). Thence, we move to show (35).

We first use the Eckhart-Young theorem and the fact that $r(XA) \leq r$ to deduce $\|Y - (PY)_r\|^2 \leq \|Y - XA\|^2 = \|E\|^2$. Hence

$$\widehat{\sigma}_r^2 \leq \frac{nm}{nm - \lambda r} \widehat{\sigma}^2 \leq \rho \widehat{\sigma}^2$$

On the other hand, Weyl's inequality shows that $d_{r+i}(PY) \geq d_{2r+i}(PE)$ for $0 \leq i \leq N - r$ by defining $d_k(PE) := 0$ for $k > N$, and we obtain

$$\begin{aligned} \widehat{\sigma}_r^2 \quad &= \quad \frac{\|Y - (PY)_r\|^2}{nm - \lambda r} \quad = \quad \frac{\|Y - PY\|^2 + \sum_{j=r+1}^N d_j^2(PY)}{nm - \lambda r} \\ &\geq \quad \frac{\|E - PE\|^2 + \sum_{j=2r+1}^N d_j^2(PE)}{nm - \lambda r} \\ &= \quad \frac{\|E - (PE)_{2r \wedge N}\|^2}{nm - \lambda r}. \end{aligned} \tag{36}$$

If $r \geq N/2$, it gives

$$\widehat{\sigma}_r^2 \geq \frac{\|E - (PE)_N\|^2}{nm - \lambda r} \geq \frac{\|E - (PE)_N\|^2}{nm - (\lambda N)/2}$$

It is worth noting that $r \geq N/2$ always guarantees a positive denominator. Otherwise, we have

$$\widehat{\sigma}_r^2 \geq \frac{\|E - (PE)_{2r}\|^2}{nm - \lambda r}.$$

We conclude the proof by invoking Lemma 14. □

Lemma 14. *On the event* $\{\lambda \geq 2d_1^2(PE)/\widehat{\sigma}^2\}$, *we have*

$$\frac{\|E\|^2}{nm} \leq \frac{\|E-(PE)_2\|^2}{nm-\lambda} \leq \frac{\|E-(PE)_4\|^2}{nm-2\lambda} \leq \cdots \leq \frac{\|E-(PE)_{(2K)\wedge N}\|^2}{nm-K\lambda}. \tag{37}$$

In particular, if λ *satisfies*

$$\lambda \geq \frac{\sum_{j=1}^{(2k)\wedge N} d_j^2(PE)}{k\widehat{\sigma}^2}$$

for any given $1 \leq k \leq r$, *then*

$$\frac{\|E\|^2}{nm} \leq \frac{\|E-(PE)_{(2k)\wedge N}\|^2}{nm-\lambda k} \leq \frac{\|E-(PE)_{(2r)\wedge N}\|^2}{nm-\lambda r} \tag{38}$$

Proof of Lemma 14. Both monotonicity claims (37) and (38) follow trivially if $2K > N$. We therefore concentrate on the case $2K \leq N$ only.

We first show (37) by using the same argument as in Propositions 1 and 2. For $k = 0, 1, \cdots, K$, we define

$$e_k := \frac{\|E-(PE)_{(2k)}\|^2}{nm-k\lambda} \tag{39}$$

and we observe that

$$\begin{aligned} e_k \leq e_{k+1} &\iff \frac{\|E-(PE)_{(2k)}\|^2}{nm-k\lambda} \leq \frac{\|E-(PE)_{(2k+2)}\|^2}{nm-(k+1)\lambda} \\ &\iff \frac{\|E-(PE)_{(2k)}\|^2}{nm-k\lambda} \leq \frac{\|E-(PE)_{(2k)}\|^2 - d_{2k+1}^2(PE) - d_{2k+2}^2(PE)}{nm-k\lambda-\lambda} \\ &\iff \frac{d_{2k+1}^2(PE)+d_{2k+2}^2(PE)}{\lambda} \leq \frac{\|E-(PE)_{(2k)}\|^2}{nm-k\lambda} = e_k. \end{aligned} \tag{40}$$

For $k = 0$, we find that $e_0 \leq e_1$ is equivalent to $d_1^2(PE) + d_2^2(PE) \leq \lambda\widehat{\sigma}^2$, which is precisely our condition on λ. Using the same reasoning as in the proof of Proposition 1, the right hand side of (40) is equivalent to

$$\frac{d_{2k+3}^2(PE)+d_{2k+4}^2(PE)}{\lambda} \leq \frac{\|E-(PE)_{(2k+2)}\|^2}{nm-(k+1)\lambda} = e_{k+1},$$

which is, by the same argument above, equivalent with $e_{k+1} \leq e_{k+2}$. Since we assumed $d_1^2(PE) + d_2^2(PE) \leq \lambda\widehat{\sigma}^2$, we conclude that $e_0 \leq e_1 \leq e_2 \leq e_3 \leq \ldots$, etcetera. This proves (37).

In fact, this proves that the sequence of $\{e_k\}$ *could only be either globally monotonic or monotone increasing at first, then monotone decreasing.*

Next we show (38). Notice that, as in the proof of Proposition 2,

$$\frac{\|E\|^2}{nm} \le \frac{\|E-(PE)_{2k}\|^2}{nm-\lambda k} \iff \lambda \ge \frac{\sum_{j=1}^{(2k)\wedge N} d_j^2(PE)}{k\widehat{\sigma}^2}.$$

Combining this equivalence with the monotone or two-sided monotone property of $\{e_k\}$ concludes the proof for $k \le r$. □

A.7 Proof of Theorem 8

We define $\mathcal{C} = \left\{(1-C_1)\sigma^2 \le \widehat{\sigma}^2 \le (1+C_1)\sigma^2\right\}$ for given $0 < C_1 < 1/2$ such that $C = 2\sqrt{1+C_1}$. Then we have

$$\begin{aligned}
&\mathbb{P}\left\{\lambda \le 2d_1^2(PE)/\widehat{\sigma}^2 \text{ or } d_s(XA) \le 2\widehat{\sigma}\sqrt{\rho\lambda}\right\} \\
&\qquad \le \mathbb{P}\left\{d_1^2(PE) \ge (1+C_0)^2(\sqrt{m}+\sqrt{q})^2\sigma^2\right\} \\
&\qquad\quad + \mathbb{P}\left\{d_s(XA) \le 2\sigma\sqrt{\rho(1+C_1)\lambda}\right\} + \mathbb{P}\{\mathcal{C}^c\} \\
&\qquad = \mathbb{P}\left\{d_1^2(PE) \ge (1+C_0)^2(\sqrt{m}+\sqrt{q})^2\sigma^2\right\} + \mathbb{P}\{\mathcal{C}^c\} \\
&\qquad \le \exp\left\{-C_0^2(m+q)/2\right\} + 2\exp\left\{-3C_1^2 nm/16\right\}
\end{aligned}$$

using (13), (14) and (15). □

Appendix B Oracle inequality

While our main interest is the study of $\widehat{k}$, we still mention the oracle inequality for our estimator of the mean XA. It improves the result stated in Giraud [18]. We first present a general oracle inequality for the case when the true rank is unknown. Then based on the overwhelming probability of the consistency of our selected $\widehat{k}$, we focus on the oracle inequality of prediction error for the given true rank. Given $\widehat{k}$ in (10), we write $\widehat{A}$ for the $p \times m$ matrix satisfying $X\widehat{A} = (PY)_{\widehat{k}}$.

Theorem 15. *Let $C > 2$. On the event $\lambda \ge Cd_1^2(PE)/\widehat{\sigma}^2$, we have*

$$\|X\widehat{A} - XA\|^2 \le \frac{C+2}{C-2} \min_{0\le k\le K} \left\{ \left(\frac{C+2}{C-2} + 8(\rho-1)\right) \sum_{j>k} d_j^2(XA) + 3\rho\lambda\widehat{\sigma}^2 k \right\}.$$

Proof. First we notice that

$$\frac{\|Y-(PY)_{\widehat{k}}\|^2}{nm-\lambda\widehat{k}} \le \frac{\|Y-(PY)_k\|^2}{nm-\lambda k} \le \frac{\|Y-(XA)_k\|^2}{nm-\lambda k}.$$

The first inequality follows from the optimality of $\widehat{k}$; the second inequality is a consequence of the Pythagoras and the Eckhart-Young theorems [14]. Rewriting the above display yields

$$\|Y - (PY)_{\widehat{k}}\|^2 \;\leq\; \|Y - (XA)_k\|^2 \left\{1 + \frac{2\lambda k}{nm - \lambda k} - \frac{\lambda(k+\widehat{k})}{nm - \lambda k}\right\},$$

and after working out the squares, we obtain

$$\begin{aligned}\|(PY)_{\widehat{k}} - XA\|^2 \;\leq\; & \sum_{j>k} d_j^2(XA) + \frac{2\lambda k}{nm - \lambda k}\|Y - (XA)_k\|^2 \\ & - \frac{\lambda(k+\widehat{k})}{nm - \lambda k}\|Y - (XA)_k\|^2 + 2\langle E, (PY)_{\widehat{k}} - (XA)_k\rangle. \end{aligned} \tag{41}$$

Since both $(XA)_k$ and $(PY)_{\widehat{k}}$ are in the column space of P, cf. Giraud [18] (page 124), we have $2\langle E, (PY)_{\widehat{k}} - (XA)_k\rangle = 2\langle PE, (PY)_{\widehat{k}} - (XA)_k\rangle$, and the norm duality and the elementary inequality $2xy \leq ax^2 + y^2/a$ for $a > 0$, yield

$$\begin{aligned} 2\langle E, X\widehat{A} - XB\rangle \leq\; & 2d_1(PE) \cdot \sqrt{k+\widehat{k}} \cdot \|(PY)_{\widehat{k}} - (XA)_k\| \\ \leq\; & (a+b)(k+\widehat{k})d_1^2(PE) + \frac{1}{a}\|(PY)_{\widehat{k}} - XA\|^2 + \frac{1}{b}\|XA - (XA)_k\|^2 \end{aligned} \tag{42}$$

for any $a, b > 0$. Moreover, we find

$$\begin{aligned}\|Y - (XA)_k\|^2 &= \|XA - (XA)_k\|^2 + \|E\|^2 + 2\langle E, XA - (XA)_k\rangle \\ &\leq 3\|XA - (XA)_k\|^2 + \frac{3}{2}\|E\|^2 \end{aligned} \tag{43}$$

and

$$\|Y - (XA)_k\|^2 \;\geq\; \frac{1}{2}\|E\|^2 - \|XA - (XA)_k\|^2 \tag{44}$$

Finally, combining (41) with (42), (43) and (44) gives

$$\begin{aligned}\frac{a-1}{a}\|X\widehat{A} - XA\|^2 \leq & \left[\frac{b+1}{b} + \frac{6\lambda k}{(nm - \lambda k)} + \frac{\lambda(k+\widehat{k})}{(nm - \lambda k)}\right]\|XA - (XA)_k\|^2 \\ & + \frac{3\lambda k}{nm - \lambda k}\|E\|^2 + (k+\widehat{k})\left[(a+b)d_1^2(PE) - \frac{\lambda/2}{nm - \lambda k}\|E\|^2\right]. \end{aligned}$$

Use $\rho \geq (nm)/(nm - \lambda\ell)$ for all $\ell \leq K$, set $a = 1 + b$ and $b = C/2$, and the claim follows after a little algebra. □

Contrary to the rank consistency results in Sections 3 and 4, no lower bound on the non-zero singular values $d_j(XA)$ of the signal XA is required. And there is no need for any assumption of X. It is clear that our selected estimator achieves the optimal bias-variance tradeoff, as discussed in [7, 17, 18]. In the above oracle inequality, the quantity $\rho = nm/(nm - \lambda K)$ could be large under some circumstances. Nevertheless, the very mild rank constraint (18) guarantees a small upper bound $\rho \le \frac{1+\delta_1}{\delta_1} \wedge \frac{1}{\delta_2}$ with any $\delta_1 > 0$ and $\delta_2 \in (0, 1)$. We will prove this claim now. For given λ, we have $K = N \wedge \lfloor nm/\lambda \rfloor$. We consider the following three cases.

(1) If $nm \ge (1 + \delta_1)\lambda N$, then $K = N$ and $0 \le r \le K = N$ and we have

$$1 \le \frac{nm}{nm - \lambda r} \le \frac{nm}{nm - \lambda N} \le \frac{(1 + \delta_1)\lambda N}{\delta_1 \lambda N} = \frac{1 + \delta_1}{\delta_1}$$

(2) If $\lambda N \le nm \le (1 + \delta_1)\lambda N$, then $K = N$ and $0 \le r \le (1 - \delta_2)N$ and we obtain

$$1 \le \frac{nm}{nm - \lambda r} \le \frac{nm}{nm - (1 - \delta_2)\lambda N} \le \frac{\lambda N}{\delta_2 \lambda N} = \frac{1}{\delta_2}$$

(3) If $qm \le nm \le \lambda N$, then $K = \lfloor nm/\lambda \rfloor$ and $0 \le r \le (1 - \delta_2)\lfloor nm/\lambda \rfloor$ and

$$1 \le \frac{nm}{nm - \lambda r} \le \frac{nm}{nm - (1 - \delta_2)nm} = \frac{1}{\delta_2}$$

as claimed.

Theorem 15 and the exponential inequalities in (13) – (15) immediately yield the following corollary.

Corollary 16. *Suppose E has i.i.d. $N(0, \sigma^2)$ entries. For $\lambda > 2(\sqrt{m} + \sqrt{q})^2$, the event*

$$\|X\widehat{A} - XA\|^2 \lesssim \rho \min_{0 \le k \le K} \left\{ \sum_{j>k} d_j^2(XA) + \lambda \sigma^2 k \right\} \tag{45}$$

holds with probability tending to 1, as $nm \to \infty$ and $q + m \to \infty$.

We use the notation $\lesssim$ for inequalities that hold up to multiplicative constants. The probability of event (45) converges (to 1) exponentially fast in mn and $m + q$.

Proof. We can write $\lambda := C'(\sqrt{m}+\sqrt{q})^2$ for some $C' = C(1+C_0)^2/(1-C_1) > 2$ with $C_0 > 0$, $0 < C_1 < 1$ and C equal to the one in Theorem 15. From inequalities (13) – (15), we have

$$\begin{aligned}\mathbb{P}\left\{Cd_1^2(PE) \geq \lambda\frac{\|E\|^2}{nm}\right\} \leq\ & \mathbb{P}\left\{d_1^2(PE) \geq (1+C_0)^2(\sqrt{m}+\sqrt{q})^2\sigma^2\right\} \\ &+\ \mathbb{P}\left\{\|E\|^2 \leq (1-C_1)nm\sigma^2\right\} \\ \leq\ & \exp\left\{-C_0^2(\sqrt{m}+\sqrt{q})^2/2\right\} + \exp\left\{-C_1^2 nm/4\right\}\end{aligned}$$

which proves the claim. □

Since the selected rank satisfies $\widehat{k} = r$ with overwhelming probability, we specialize to the case when r is known. Theorem 17 provides a cleaner and tighter bound for the prediction error.

Theorem 17. *On the event (24) and (25) in Theorem 10, we have*

$$\|X\widehat{A} - XA\|^2 \leq 4rd_1^2(PE)$$

Proof. Theorem 10 implies that $\widehat{k} := \widehat{k}_s = r$. From the optimality of $X\widehat{A}$, we have

$$\|Y - X\widehat{A}\|^2 \leq \|Y - (XA)_r\|^2 = \|E\|^2$$

on the event $\{\widehat{r} = r\}$. This implies

$$\|X\widehat{A} - XA\|^2 \leq 2|\langle E, X\widehat{A} - (XA)_r\rangle| \leq 2\|X\widehat{A} - XA\|(\sqrt{r}d_1(PE))$$

and the result follows. □

Corollary 18. *Assume E has i.i.d. $N(0,\sigma^2)$ entries. On the event (29) in Theorem 13, if we update λ according to (26), then we have $\widehat{k} = r$ with overwhelming probability. Hence, for some constant $C > 4$, we have*

$$\mathbb{P}\left\{\|X\widehat{A} - XA\|^2 \leq Cr(\sqrt{q}+\sqrt{m})^2\sigma^2\right\} \to 1$$

as $nm \to \infty$ and $q + m \to \infty$.

The convergence rate in Corollary 18 is exponentially fast. The proof becomes the same as that of Theorem 13.

Appendix C Proofs of Section 4

C.1 Proof of Proposition 9

First, we show $\widehat{k}_t \leq r$ for all $0 \leq t \leq T$. From Corollary 7 and our procedure, we know $1 \leq \widehat{k}_0 \leq r$. Let $\widehat{\lambda}_1$ chosen from $\widehat{k}_0$ using updating rule (22). In order to guarantee $\widehat{k}_1 \leq r$, using the bounds (36) and (19), we need to show

$$d_1^2(PE) \leq \frac{\|E - (PE)_{(2r)\wedge N}\|^2}{nm/\widehat{\lambda}_1 - r}. \tag{46}$$

The choice

$$\widehat{\lambda}_1 \geq \frac{nm\left[d^2_{2\widehat{k}_1-1}(PE) + d^2_{2\widehat{k}_1}(PE)\right]}{\|E - (PE)_{(2\widehat{k}_1)\wedge N}\|^2 + \widehat{k}_1\left[d^2_{2\widehat{k}_1-1}(PE) + d^2_{2\widehat{k}_1}(PE)\right]}$$

and the monotonicity of (37) in the proof of Proposition 14 guarantee

$$\frac{\|E - (PE)_{(2\widehat{k}_1-2)\wedge N}\|^2}{nm - \widehat{\lambda}_1(\widehat{k}_1 - 1)} \leq \frac{\|E - (PE)_{(2\widehat{k}_1)\wedge N}\|^2}{nm - \widehat{\lambda}_1\widehat{k}_1} \leq \frac{\|E - (PE)_{(2r)\wedge N}\|^2}{nm - \widehat{\lambda}_1 r} \tag{47}$$

provided that $\widehat{k}_1 \leq r$. On the other hand,

$$\widehat{\lambda}_1 \geq \frac{nm}{\|E - (PE)_{(2\widehat{k}_1)\wedge N}\|^2/d_1^2(PE) + \widehat{k}_1} \iff d_1^2(PE) \leq \frac{\|E - (PE)_{(2\widehat{k}_1)\wedge N}\|^2}{nm/\widehat{\lambda}_1 - \widehat{k}_1} \tag{48}$$

Combining (47) with (48) proves (46). After repeating this argument, we find

$$d_1^2(PE) \leq \frac{\|E - (PE)_{(2r)\wedge N}\|^2}{nm/\widehat{\lambda}_t - r} \leq \frac{\|Y - (PY)_r\|^2}{nm/\widehat{\lambda}_t - r} \tag{49}$$

for any $t \geq 1$, which implies $\widehat{k}_t \leq r$. Note that, by using the same arguments, the following holds

$$d_1^2(PE) \leq \frac{\|Y - (PY)_r\|^2}{nm/\widehat{\lambda}_{k_t} - r} \tag{50}$$

for all $\widehat{\lambda}_{k_t}$ which is computed from (22) by using any $1 \leq k_t \leq r$ instead of $\widehat{k}_t$.

Next we show (22) is decreasing in $\widehat{k}_t$. To simplify notation, we write k for $\widehat{k}_t$. Since it is immediate to verify the decreasing property when $2k \geq N$, we consider $1 \leq k < N/2$ only. First, we claim

$$\frac{nm}{\|E - (PE)_{2k}\|^2 / \left[d^2_{2k-1}(PE) + d^2_{2k}(PE)\right] + k}$$

is decreasing in k. It suffices to show

$$a_k := \frac{\|E - (PE)_{2k}\|^2}{d_{2k-1}^2(PE) + d_{2k}^2(PE)} + k \tag{51}$$

is increasing. This is true by noting

$$a_k - a_{k-1} = \frac{\|E - (PE)_{2k-2}\|^2}{d_{2k-1}^2(PE) + d_{2k}^2(PE)} - \frac{\|E - (PE)_{2k-2}\|^2}{d_{2k-3}^2(PE) + d_{2k-2}^2(PE)} > 0$$

Then we only need to show, if $d_1^2(PE) \geq d_{2k-1}^2(PE) + d_{2k}^2(PE)$ that

$$b_k := \frac{nm}{\|E - (PE)_{2k}\|^2/d_1^2(PE) + k} \tag{52}$$

is decreasing in k, which is implied by the fact that b_k starts decreasing at k when $d_1^2(PE) \geq d_{2k+1}^2(PE) + d_{2k+2}^2(PE)$. This proves $\widehat{\lambda}_{t+1}$ is decreasing in $\widehat{k}_t$.
Finally, we show that, for given $\widehat{\lambda}_t > \widehat{\lambda}_{t+1}$, we have $\widehat{k}_t \leq \widehat{k}_{t+1}$. Suppose $\widehat{k}_t > \widehat{k}_{t+1}$, we obtain

$$\frac{\|Y - (PY)_{\widehat{k}_{t+1}}\|^2}{nm - \widehat{\lambda}_{t+1}\widehat{k}_{t+1}} \leq \frac{\|Y - (PY)_{\widehat{k}_t}\|^2}{nm - \widehat{\lambda}_{t+1}\widehat{k}_t} \iff \frac{\sum_{j=\widehat{k}_{t+1}}^{\widehat{k}_t} d_1^2(PE)}{\widehat{k}_t - \widehat{k}_{t+1}} \leq \frac{\|Y - (PY)_{\widehat{k}_t}\|^2}{nm/\widehat{\lambda}_{t+1} - \widehat{k}_t}$$

Similarly,

$$\frac{\|Y - (PY)_{\widehat{k}_{t+1}}\|^2}{nm - \widehat{\lambda}_t\widehat{k}_{t+1}} \geq \frac{\|Y - (PY)_{\widehat{k}_t}\|^2}{nm - \widehat{\lambda}_t\widehat{k}_t} \iff \frac{\sum_{j=\widehat{k}_{t+1}}^{\widehat{k}_t} d_1^2(PE)}{\widehat{k}_t - \widehat{k}_{t+1}} \geq \frac{\|Y - (PY)_{\widehat{k}_t}\|^2}{nm/\widehat{\lambda}_t - \widehat{k}_1}$$

which is a contradiction as $\widehat{\lambda}_t > \widehat{\lambda}_{t+1}$. This concludes the proof. □

C.2 Proof of Theorem 10

Our assumption

$$d_{k_0}(XA) \geq d_1(PE) + \frac{\|Y - (PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{k_{-1}} - r}},$$

implies that $\widehat{k}_0 \in [k_0, r]$ by Theorem 6. This, in turn, from the decreasing property of $\widehat{\lambda}_t$ in Proposition 9, yields $\widehat{\lambda}_{\widehat{k}_0} \leq \widehat{\lambda}_{k_0}$, where $\widehat{\lambda}_{\widehat{k}_0}$ is computed from (22) by using $\widehat{k}_0$. Hence

$$d_{k_1}(XA) \;\geq\; d_1(PE) + \frac{\|Y - (PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{k_0} - r}} \;\geq\; d_1(PE) + \frac{\|Y - (PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{\widehat{k}_0} - r}}$$

and by the same argument above, we have $\widehat{k}_1 \in [k_1, r]$. Then the proof follows by repeating these arguments and the fact that $k_T = r$. □

C.3 Proof of Proposition 11

First, we focus on $\widehat{\lambda}_0$ and $\widetilde{\lambda}_0$. From (13) and (15), for any $0 < \varepsilon < 1/2$, we have

$$\begin{aligned}\mathbb{P}\left\{\widetilde{\lambda}_0 \geq \widehat{\lambda}_0\right\} \geq\ & \mathbb{P}\left\{(1+\varepsilon)(\sqrt{m}+\sqrt{q})^2\sigma^2 \geq d_1^2(PE)\right\} + \mathbb{P}\left\{\widehat{\sigma}^2 \geq (1+\varepsilon)\sigma^2\right\} \\ \geq\ & 1 - \exp(-\varepsilon(\sqrt{m}+\sqrt{q})^2/2) - \exp(-3nm\varepsilon^2/16) \end{aligned} \tag{53}$$

For $t \geq 0$, from (22) and (26), it is equivalent to show

$$\mathbb{P}\left\{\frac{\|E-(PE)_{2k_t}\|^2}{\widehat{U}_t} \geq \frac{(1-\varepsilon)nm - \widetilde{R}_t}{\widetilde{U}_t}\right\} \to 1, \quad 2k_t < N$$

and

$$\mathbb{P}\left\{\frac{\|E-PE\|^2}{d_1^2(PE)} \geq \frac{(1-\varepsilon)(n-q)m}{(\sqrt{q}+\sqrt{m})^2}\right\} \to 1, \quad 2k_t \geq N$$

where $\widetilde{U}_t$ and $\widetilde{R}_t$ are defined in (27) and (28) but based on k_t and $\widehat{U}_t$ replaces $\widehat{k}_t$ by k_t.
To simplify notation in the following, we write k for k_t. First we consider $2k < N$ and aim to bound $\|(PE)_{2k}\|^2 = \sum_{j=1}^{2k} d_j^2(PE)$. Let $U^T\Lambda U$ be the eigenvalue decomposition of the projection matrix P. Then $d_j^2(PE) = \lambda_j(E^TPE) = d_j^2(\Lambda U^T E)$. The latter one has $q \times m$ non-zero Gaussian entries, see Lemma 3 in Bunea, She and Wegkamp [7]. From the Gaussian concentration inequality of Giraud [18] (page 221), by noting that $\|(X)_{2k}\|$ is a $2k$-Lipschitz function, we have

$$\mathbb{P}\left\{\|(PE)_{2k}\| \leq E[\|(PE)\|_{2k}\|] + \sigma t\right\} \geq 1 - \exp\left\{-t^2/(8k^2)\right\},\ \forall\ t > 0.$$

From $\mathbb{E}[\|(PE)_{2k}\|] \leq \sqrt{\mathbb{E}[\|(PE)_{2k}\|^2]}$ and the upper bound of $\mathbb{E}[\|G\|_{2,r}^2]$ in Giraud [17] (Fact 1 on page 778), we obtain

$$\mathbb{P}\left\{\|(PE)_{2k}\|^2 \leq (1+t')^2 R_k\sigma^2\right\} \geq 1 - \exp\left\{-t'^2 R_k/(8k^2)\right\}$$

for all $t' > 0$ and $2k < N$. Similarly, since each singular value operator is a 1-Lipschitz function, the same Gaussian concentration inequality gives us

$$\mathbb{P}\left\{d_j^2(PE) \leq (1+t)^2(\sqrt{M}+\sqrt{N-j+1})^2\sigma^2\right\} \geq 1 - \exp\left\{-t^2\left(\sqrt{M}+\sqrt{N-j+1}\right)^2/2\right\}$$

for all $t > 0$ and $j = 1, \ldots, N$. We denote, for some $C_0 > 0$ and $0 < C_1 < 1$,

$$\begin{aligned}
\mathcal{A} &= \left\{d_1^2(PE) \le (1+C_0)^2(\sqrt{m}+\sqrt{q})^2\sigma^2\right\}, \\
\mathcal{B} &= \left\{\widehat{\sigma}^2 \ge (1-C_1)nm\sigma^2\right\}, \\
\mathcal{C} &= \left\{\|(PE)_{2k}\|^2 \le (1+C_0)^2 R_k \sigma^2\right\}, \\
\mathcal{D} &= \{d_{2k-1}^2(PE) \le (1+C_0)^2(\sqrt{M}+\sqrt{N-2k+2})^2\sigma^2\}, \\
\mathcal{E} &= \{d_{2k}^2(PE) \le (1+C_0)^2(\sqrt{M}+\sqrt{N-2k+1})^2\sigma^2\},
\end{aligned}$$

Invoking (13), (14) and (15) and choosing $1-\varepsilon = (1-C_1)/(1+C_0)^2$, we obtain

$$\begin{aligned}
&\mathbb{P}\left\{\frac{\|E-(PE)_{2k}\|^2}{U_k} \ge \frac{(1-\varepsilon_1)nm - R_k}{\widetilde{U}_k}\right\} \\
&\quad\ge \mathbb{P}\{\mathcal{A}\cap\mathcal{B}\cap\mathcal{C}\cap\mathcal{D}\cap\mathcal{E}\} \\
&\quad\ge 1 - \exp\left\{-C_1^2 nm/4\right\} - \exp\left\{-C_0^2(\sqrt{m}+\sqrt{q})^2/2\right\} \\
&\qquad - \exp\left\{-C_0^2 R_k/(8k^2)\right\} - \exp\left\{-C_0^2\left(\sqrt{M}+\sqrt{N-2k+2}\right)^2/2\right\} \\
&\qquad - \exp\left\{-C_0^2\left(\sqrt{M}+\sqrt{N-2k+1}\right)^2/2\right\}.
\end{aligned}$$

Finally, if $2k \ge N$, we observe that $\|E-PE\|^2$ has a central $\sigma^2\chi^2_{(n-q)m}$ distribution. Denote $\mathcal{B}' = \{\|E-PE\|^2 \ge (1-C_1)(n-q)m\sigma^2\}$ for $0 < C_1 < 1$. By using the same ε as above, we have

$$\begin{aligned}
\mathbb{P}\left\{\frac{\|E-PE\|^2}{d_1^2(PE)} \ge \frac{(1-\varepsilon)(n-q)m}{(\sqrt{q}+\sqrt{m})^2}\right\} \quad &\ge \quad 1 - \exp\left\{-3C_1^2(n-q)m/16\right\} \\
&\quad - \exp\left\{-C_0^2(\sqrt{m}+\sqrt{q})^2/2\right\},
\end{aligned}$$

which completes the proof. □

C.4 Proof of Proposition 12

For clarification, here we use the explicit expression $\widetilde{\lambda}_{\widetilde{k}_t}$. To show $\widetilde{\lambda}_{\widetilde{k}_t}$ is decreasing in $\widetilde{k}_t$, for simplicity of notation, we show the decreasing property of $\widetilde{\lambda}_k$ obtained from (26) by using k in lieu of $\widetilde{k}_t$. Correspondingly, $\widetilde{U}_t$ and $\widetilde{R}_t$ are also calculated based on k in place of $\widetilde{k}_t$.

Note that $\widetilde{\lambda}_1 \le \widetilde{\lambda}_0$ is straightforward. Since the decreasing property of (26) is easily to see when $2k \ge N$, we focus on the case $2k < N$. Observe that $\widetilde{\lambda}_k$ is decreasing when $\widetilde{R}_t = qm - \sum_{j=1}^{2k}(\sqrt{M}-\sqrt{j})^2$. Therefore, we need to show $\widetilde{\lambda}_k$ is decreasing in k when

$\widetilde{R}_t = \sum_{j=1}^{2k}(\sqrt{M}+\sqrt{N-j+1})^2$. It suffices to prove the increasing property of

$$[(1-\varepsilon)nm - R_t]/\widetilde{U}_t + k = \left[(1-\varepsilon)nm - \sum_{j=1}^{2k}(\sqrt{M}+\sqrt{N-j+1})^2\right]/\widetilde{U}_t + k$$

where

$$\widetilde{U}_t := \max\left\{(\sqrt{m}+\sqrt{q})^2, \left(\sqrt{M}+\sqrt{N-2k+2}\right)^2 + \left(\sqrt{M}+\sqrt{N-2k+1}\right)^2\right\}$$

When $\widetilde{U}_t = \left(\sqrt{M}+\sqrt{N-2k+2}\right)^2 + \left(\sqrt{M}+\sqrt{N-2k+1}\right)^2$, this becomes the same form as (51) in the proof of Proposition 9. On the other hand, when $\widetilde{U}_t = (\sqrt{m}+\sqrt{q})^2$, this is the similar to (52). Therefore, the same argument can be used to prove that $\widetilde{\lambda}_k$ is decreasing. We proceed to prove $\mathbb{P}\{\widetilde{k}_t \le r\} \to 1$. From our starting point (19) with $\widetilde{\lambda}_t \equiv \widetilde{\lambda}_{\widetilde{k}_{t-1}}$ in mind, we need to show

$$\mathbb{P}\left\{d_1(PE) \le \frac{\|Y-(PY)_r\|}{\sqrt{nm/\widetilde{\lambda}_{\widetilde{k}_{t-1}} - r}}\right\} \to 1$$

First, Proposition 11 gives $\mathbb{P}\{\widetilde{\lambda}_{\widetilde{k}_{t-1}} \ge \widehat{\lambda}_{\widetilde{k}_{t-1}}\} \to 1$ for any $t \ge 0$. On the other hand, from (50) in the proof of Proposition 9, the following holds

$$d_1(PE) \le \frac{\|Y-(PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{k_t} - r}}, \qquad \text{for any } 1 \le k_t \le r$$

Combining together concludes the proof. □

C.5 Proof of Theorem 13

First, we notice that, from (53),

$$\mathbb{P}\left\{2d_1^2(PE)/\widehat{\sigma}^2 \le \widetilde{\lambda}_0\right\} \ge 1 - \exp\left(-\varepsilon(\sqrt{m}+\sqrt{q})^2/2\right) - \exp\left(-3\varepsilon^2 nm/16\right)$$

for any $0 < \varepsilon < 1/2$. In addition, from (50) and Proposition 12, we have

$$\mathbb{P}\left\{d_1(PE) \le \frac{\|Y-(PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{k_{t-1}} - r}} \le \frac{\|Y-(PY)_r\|}{\sqrt{nm/\widetilde{\lambda}_{k_{t-1}} - r}}\right\} \to 1.$$

This further implies that the probability of event

$$\mathcal{A} := \left\{d_{k_t}(XA) \ge 2\widehat{\sigma}\sqrt{\widetilde{\rho}_t\widetilde{\lambda}_{k_{t-1}}} \overset{(i)}{\ge} d_1(PE) + \frac{\|Y-(PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{k_{t-1}} - r}}\right\}$$

goes to 1, where (i) comes from the fact that $\|E\| \geq \|Y - (PY)_r\|$. Therefore, by choosing $C = 2(1+\varepsilon)$, we have

$$\mathbb{P}\left\{ d_{k_t}(XA) \geq C\sigma\sqrt{\widetilde{\rho}_t \widetilde{\lambda}_{k_{t-1}}} \geq d_1(PE) + \frac{\|Y - (PY)_r\|}{\sqrt{nm/\widehat{\lambda}_{k_{t-1}} - r}} \right\} \geq \mathbb{P}(\mathcal{A}) - \exp(-3nm\varepsilon^2/16)$$

for $t = 0, 1, \ldots, T$ and any $0 < \varepsilon < 1/2$. Invoke Theorem 10 to conclude the proof. □

References

[1] Anderson, T.W. (1951). Estimating linear restrictions on regression coefficients for multivariate normal distributions. *Annals of Mathematical Statistics*, **22**, 327-351.

[2] Anderson, T.W. (1999). Asymptotic distribution of the reduced rank regression estimator under general conditions. *Annals of Statistics*, **27**(4), 1141 - 1154.

[3] Anderson, T.W. (2002). Specification and misspecification in reduced rank regression. *Sankya*, **64**, Series A, 193 - 205.

[4] Aldrin, M. (1996). Moderate projection pursuit regression for multivariate response data. *Comput. Statist. Data Anal*, **21**, 501–531.

[5] Birgé, L. and Massart, P. (2007). Minimal penalties for Gaussian model selection. *Probability Theory and Related Fields*. **138**, no 1-2, 33–73.

[6] Bunea, F., Tsybakov, A.B. and Wegkamp, M.H. (2007). Aggregation for Gaussian regression. *Annals of Statistics*, **35**(4), 1674–1697.

[7] Bunea, F., She, Y. and Wegkamp, M.H. (2011). Optimal Selection of Reduced Rank Estimators of High-Dimensional Matrices. *Annals of Statistics*, **39**, 1282-1309.

[8] Bunea, F., She, Y. and Wegkamp, M.H. (2012). Joint variable and rank selection for parsimonious estimation of high-dimensional matrices. *Annals of Statistics*, **40**, no. 5, 2359–2388.

[9] Bhlmann, P. and Van de Geer, S. (2011). *Statistics for High-Dimensional Data* Springer.

[10] Candès, E.J. and Tao, T. (2009). The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inform. Theory*, **56**(5), 2053-2080.

[11] Candès, E.J. and Plan, Y. (2010). Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements. *arXiv:1001.0339 [cs.IT]*

[12] Cavalier, L., Golubev, G.K., Picard, D. and Tsybakov, A.B. (2002). Oracle inequalities for inverse problems. *Annals of Statistics*, **30**: 843–874.

[13] Chafaï, D. (2009). Singular values of random matrices. *Lecture Notes*. Available at authors website.

[14] Eckart, C. and Young, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika*. **1**: 211. doi:10.1007/BF02288367.

[15] Fazel, M. (2002). Matrix rank minimization with applications. PhD thesis, Stanford University.

[16] Horn, R.A. and Johnson, C.R. (2013). Matrix analysis. Second edition. *Cambridge University Press, Cambridge*. xviii+643 pp. ISBN: 978-0-521-54823-6

[17] Giraud, C. (2011). Low rank multivariate regression. *Electronic Journal of Statistics*, 5, 775-799.

[18] Giraud, C. (2015). Introduction to high-dimensional statistics. *Monographs on Statistics and Applied Probability*, **139**. CRC Press, Boca Raton, FL. xvi+255 pp. ISBN: 978-1-4822-3794-8.

[19] Izenman, A.J.(1975). Reduced-Rank Regression for the Multivariate Linear Model. *Journal of Multivariate Analysis*, **5**, 248–262.

[20] Izenman, A.J. (2008). *Modern Multivariate. Statistical Techniques: Regression, Classification and Manifold Learning.* Springer, New York.

[21] Johnstone, I. M. (2001). Chi-square oracle inequalities, in M. de Gunst, C. Klaassen & A. van der Vaart, eds, *Festschrift for Willem R. van Zwet*. Volume 36 of IMS Lecture Notes - Monographs, Institute of Mathematical Statistics, 399?418.

[22] Kolmogorov, A.N. and Tikhomirov, V.M.(1961). ε-entropy and ε-capacity of sets in functions spaces. *Amer. Math. Soc. Transl.*, **17**, 277–364.

[23] Lounici, K., Pontil, M., Tsybakov, A.B. and Van de Geer, S.A. (2011). Oracle Inequalities and Optimal Inference under Group Sparsity. *Annals of Statistics*, **39**, 2164-2204.

[24] Marčenko, V.A. and Pastur, L.A. (1967). Distribution of eigenvalues in certain sets of random matrices. (Russian) *Mat. Sb. (N.S.)*, **72** (114), 507–536.

[25] Lu, Z.; Monteiro, R.D.C. and Yuan, M. (2012). Convex optimization methods for dimension reduction and coefficient estimation in multivariate linear regression. *Math. Program.* **131**, no. 1-2, Ser. A, 163–194.

[26] Negahban, S. and Wainwright, M.J. (2011). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Ann. Statist.* **39**, no. 2, 1069–1097.

[27] Obozinski, G., Wainwright, M. and Jordan, M. (2011). Support union recovery in high-dimensional multivariate regression. *Ann. Statist.* **39**, 1–47.

[28] Rao, C.R. (1978). Matrix Approximations and Reduction of Dimensionality in Multivariate Statistical Analysis. *Multivariate Analysis V.* Proceedings of the fifth international symposium of multivariate analysis; P. R. Krishnaiah Editor, North-Holland Publishing.

[29] Reinsel, G.C. and Velu, R.P. (1998). *Multivariate Reduced-Rank Regression: Theory and Applications.* Lecture Notes in Statistics, Springer, New York.

[30] Robinson, P.M. (1973). Generalized canonical analysis for time series. *Journal of Multivariate Analysis*, **3**, 141–160.

[31] Robinson, P.M. (1974). Identification, estimation and large sample theory for regression containing unobservable variables. *International Economic Review*, **15**, 680-692.

[32] Rohde, A. and Tsybakov, A.B. (2011). Estimation of high-dimensional low-rank matrices. *Ann. Statist.* **39**, no. 2, 887–930.

[33] Rudelson, M. and Vershynin, R. (2010). Non-asymptotic theory of random matrices: extreme singular values. *Proceedings of the International Congress of Mathematicians*. Volume III, 1576–1602, Hindustan Book Agency, New Delhi.

[34] Schmidt, E. (1907). Zur Theorie der linearen und nichtlinearen Integralgleichungen. (German) *Math. Ann.* **63**, no. 4, 433–476.

[35] Stewart, G.W. (1993). On the early history of the singular value decomposition. *SIAM REVIEW*, **35** 551–566.

[36] Takane, Y. and Hunter, M.A. (2001). Constrained principal component analysis: A comprehensive theory. *Applicable Algebra in Engineering, Communication, and Computing*, **12**, 391-419.

[37] Takane, Y. and Hwang, H. (2007). Regularized linear and kernel redundancy analysis. *Computational Statistics and Data Analysis*, **52**, 394-405.

[38] van der Vaart, A.W. and Wellner, J.A. (1996). *Weak convergence and Empirical Processes.* Springer-Verlag, New York.

[39] Weyl, H. (1912). Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Mathematische Annalen.* **71**, no. 4: 441-479.

[40] Yuan, M., Ekici, A., Lu, Z. and Monteiro, R. (2007). Dimension reduction and coefficient estimation in multivariate linear regression. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **69**, no. 3, 329–346.

[41] Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *J. Roy. Statist. Soc. Ser. B* **68**, 49–67.